LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

UCRL-SR-229520

Tri-Laboratory Linux Capacity Cluster 2007 SOW

Mark Seager

March 28, 2007

Disclaimer

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

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

UCRL-MI-TBD

University of California
Lawrence Livermore National Laboratory

Tri-Laboratory Linux Capacity Cluster 2007

Draft
Statement of Work
B563016
Attachment 2

Version 4
March 17, 2007

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

This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Contents

1 BACKGROUND 7

1.1 ADVANCED SIMULATION AND COMPUTING PROGRAM 7
1.2 ASC CAPACITY SYSTEMS STRATEGY 7
1.3 TRI-LABORATORY SIMULATION ENVIRONMENTS 8
1.3.1 The LLNL Simulation Environment 8
1.3.2 The LANL Simulation Environment 9
1.3.3 The SNL Simulation Environment 10
1.4 UTILIZATION OF EXISTING FACILITIES 11
1.4.1 The LLNL Existing Facilities 11
1.4.2 The LANL Existing Facilities 13
1.4.3 The Sandia Existing Facilities 14

2 TLCC SCALABLE UNIT STRATEGY AND ARCHITECTURE 18

2.1 TLCC STRATEGY 18
2.2 TLCC CLUSTER ARCHITECTURE 19
2.3 TLCC SOFTWARE ENVIRONMENT 23
2.3.1 Tri-Laboratory Build and Accept Software Stack (BASS) 23
2.3.2 LLNL Clustered High Availability Operating System (CHAOS) 24
2.3.2.1 LLNL Cluster Tools 26
2.3.2.2 Simple Linux Utility for Resource Management 26
2.3.2.3 Moab Cluster Suite 27
2.3.2.4 Lustre Cluster Wide File System 27
2.3.2.5 The Livermore Computing Linux Cluster Support Model 27
2.3.2.6 Integration Testing 28
2.3.3 LANL Linux Capacity Environment (LLCE) 28
2.3.3.1 Panasas PanFS Multi-Cluster Environment Wide Global Parallel OBSD Based File System 31
2.3.3.2 Regression analysis and Testing at LANL 31
2.3.3.3 Configuration Control 32
2.3.3.4 Parallel Programming Models & Run-Time Systems 32
2.3.3.5 Parallel I/O Libraries 32
2.3.3.6 Compilation/Build Environment 33
2.3.3.7 Debugging and Correctness Testing 33
2.3.3.8 Performance Measurement, Analysis, and Tuning 35
2.3.4 SNL Linux Software (Chimera) 36
2.3.4.1 SNL Cluster Tools 37
2.3.4.2 TORQUE Resource Manager and Job Scheduling 38
2.3.4.3 Lustre Cluster Wide File System 38
2.3.4.4 The Sandia Computing Linux Cluster Support Model 38
2.3.4.5 Integration Testing 39

3 TLCC TECHNICAL REQUIREMENTS 40

3.1 HIGH-LEVEL HARDWARE SUMMARY (TR-1) 41
3.1.1 SU High-Level Architecture 41
3.1.2 SU Requirements Summary Matrix 41
3.1.3 SU Evolution Overview (TR-1) 43
3.2 SU HARDWARE REQUIREMENTS 45
3.2.1 TLCC SMP Scalable Unit (MR) 45
3.2.2 TLCC SMP Four Socket Configuration (TR-1) 45
3.2.3 SU Peak (TR-1) 45
3.2.4 Number of SU (MR) 45
3.2.5 TLCC Node Requirements 45
3.2.5.1 Processor and Cache (TR-1) 45
3.2.5.2 Chipset and Memory Interface (TR-2) 46
3.2.5.3 Node Delivered Memory Performance (TR-1) 46
3.2.5.4 Node I/O Configuration (TR-2) 46
3.2.5.5 Node Memory (TR-1) 46
3.2.5.6 Node Power (TR-2) 46

3.2.5.7 Linux Access to Memory Error Detection (TR-1) 46
3.2.5.8 No Local Hard Disk (TR-1) 47
3.2.5.9 Node Form Factor (TR-2) 47
3.2.5.10 Integrated Management Ethernet (TR-1) 47
3.2.5.11 Node BIOS (MR) 47
3.2.5.11.1 Node BIOS Type Options (MR) 47
3.2.5.11.1.1 LinuxBIOS (TR-1) 47
3.2.5.11.1.2 Standard BIOS (TR-2) 47
3.2.5.11.2 Remote Network Boot (MR) 48
3.2.5.11.3 Remote IBA Network Boot (TR-1) 48
3.2.5.11.4 Node Initialization (TR-1) 48
3.2.5.11.5 Error/Interrupt Handling (MR) 48
3.2.5.11.6 BIOS Security (MR) 48
3.2.5.11.7 Plans and Process for Needed BIOS Updates (MR) 48
3.2.5.11.8 Failsafe/Fallback Boot Image (TR-1) 49
3.2.5.11.9 Failsafe CMOS Parameters (TR-1) 49
3.2.5.11.10 Serial Console after Failsafe/Fallback Booting (TR-1) 49
3.2.5.11.11 BIOS Upgrade and Restore (TR-1) 49
3.2.5.11.12 CMOS Parameter Manipulation (TR-1) 49
3.2.5.11.13 BIOS Command Line Interface (TR-2) 49
3.2.5.11.14 Serial Console over Ethernet Support (TR-2) 50
3.2.5.11.15 Power-On Self Test (TR-2) 50
3.2.5.11.16 BIOS Security Verification (MR) 50
3.2.5.11.16.1 BIOS Buildable Source Code (TR-1) 50
3.2.5.11.16.2 BIOS Security Certification (TR-1) 50
3.2.5.12 Programmable LED(s) (TR-3) 50
3.2.6 IBA Interconnect (MR) 50
3.2.6.1 Node HCA Functionality (TR-1) 50
3.2.6.2 Node Bandwidth, Latency and Throughput (TR-1) 51
3.2.6.3 Fully Functional IBA Topology (TR-1) 51
3.2.6.4 IBA Interconnect Reliability (TR-1) 51
3.2.6.5 Multi-SU Spine Switches (TR-1) 52
3.2.7 Remote Manageability (TR-1) 52
3.2.7.1 Traditional Remote Management Solution (TR-1) 52
3.2.7.1.1 Remote Console (TR-3) 52
3.2.7.1.2 Remote Node Power Control over Management LAN (TR-3) 53
3.2.7.2 IPMI and BMC Remote Management Solution (TR-1) 53
3.2.7.2.1 ConMan Access to Console via Serial over LAN (TR-1) 54
3.2.7.2.2 LAN PowerMan Access (TR-1) 54
3.2.7.2.3 LAN Management Access (TR-1) 54
3.2.7.2.4 IPMI User Role Division (TR-2) 54
3.2.7.2.5 Traditional Remote Management Backup Plan (TR-2) 54
3.2.7.2.6 Additional IPMI Security Requirements (MR or TR-1?) 54
3.2.7.2.6.1 Disabling a User After Failed Out-of-band Authentication Attempts (TR-1) 55
3.2.7.2.6.2 Failed IPMI Out-of-band Authentication Attempts (TR-1) 55
3.2.7.3 Remote Management Solution Requirements (TR-1) 55
3.2.7.3.1 Serial Console Redirection (TR-1) 55
3.2.7.3.2 Dedicated Serial Console Communications (TR-2) 55
3.2.7.3.3 Serial Console Efficiency (TR-1) 55
3.2.7.3.4 Flow Control (TR-1) 56
3.2.7.3.5 Peripheral Device Firmware (TR-2) 56
3.2.7.3.6 Remote Network Boot Mechanism (TR-1) 56
3.2.8 Gateway Node Requirements (TR-1) 56
3.2.8.1 Gateway Node Count (TR-1) 56
3.2.8.2 Gateway Node Configuration (TR-1) 56
3.2.8.3 Gateway Node I/O Configuration (TR-1) 56
3.2.8.4 Gateway Node 10 Gb/s Ethernet Card (TR-1) 57
3.2.8.5 Gateway Node Delivered Performance (TR-2) 57
3.2.9 Login/Service/Master Node Requirements 57
3.2.9.1 LSM Node Count (TR-1) 57
3.2.9.2 LSM Node I/O Configuration (TR-1) 57
3.2.9.3 LSM Node Ethernet Configuration (TR-1) 57
3.2.9.4 LSM Node Accessory Configuration (TR-2) 57
3.2.9.5 LSM Node Delivered Performance (TR-2) 58

3.2.10 Remote Partition Service Node Requirements 58
3.2.10.1 RPS Node Count (TR-1) 58
3.2.10.2 RPS Node I/O Configuration (TR-1) 58
3.2.10.3 RPS Node Ethernet Configuration (TR-1) 58
3.2.10.4 RPS Node RAID5 Configuration (TR-1) 58
3.2.11 SU Management Ethernet (TR-1) 59
3.2.12 SU Racks and Packaging (TR-1) 59
3.2.12.1 SU Design Optimization (TR-2) 59
3.2.12.2 Rack Height and Weight (TR-1) 59
3.2.12.3 Rack Structural Integrity (TR-2) 60
3.2.12.4 Rack Air Flow and Cooling (TR-1) 60
3.2.12.5 Rack Doors (TR-2) 60
3.2.12.6 Rack Cable Management (TR-2) 60
3.2.12.7 Rack Color (TR-3) 60
3.2.12.8 Rack Power and Cooling (TR-1) 60
3.2.12.9 Rack PDU (TR-1) 61
3.2.13 UL Listing (TR-1) 61
3.3 TLCC SOFTWARE REQUIREMENTS (TR-1) 61
3.3.1 Minimum IBA Software Stack (MR) 61
3.3.1.1 OpenFabrics Enterprise Distribution (MR) 61
3.3.2 IBA Upper Layer Protocols (TR-1) 62
3.3.3 TLCC Peripheral Device Drivers (TR-1) 62
3.3.4 RPS Node Software (TR-1) 62
3.3.5 Kernel Support for Memory Error Detection (TR-1) 63
3.3.6 Linux Access to Motherboard Sensors (TR-1) 63
3.3.7 Remote Management Software (TR-2) 63
3.3.7.1 TRMS Software (TR-1) 63
3.3.7.2 IPMI and BMC Remote Management Software (TR-1) 63
3.3.7.3 Linux Tool for LinuxBIOS Setup (TR-1) 64
3.3.7.4 Linux Tool for BIOS Upgrade (TR-1) 64
3.3.8 System Diagnostics (TR-2) 64
3.4 SU HARDWARE EVOLUTION (TR-1) 64
3.5 SU SOFTWARE EVOLUTION (TR-1) 64

4 RELIABILITY, AVAILABILITY, SERVICEABILITY (RAS) AND MAINTENANCE 65

4.1 HIGHLY RELIABLE MANAGEMENT NETWORK (TR-1) 66
4.2 TLCC NODE RELIABILITY AND MONITORING (TR-2) 66
4.3 IN PLACE NODE SERVICE (TR-1) 66
4.4 COMPONENT LABELING (TR-1) 66
4.5 FIELD REPLACEABLE UNIT (FRU) DIAGNOSTICS (TR-2) 67
4.5.1 Node Diagnostics Suite (TR-1) 67
4.5.2 Memory Diagnostics (TR-1) 67
4.5.3 IBA Diagnostics (TR-1) 67
4.5.4 Peripheral Component Diagnostics (TR-2) 68
4.6 SCALABLE SYSTEM MONITORING (TR-2) 68
4.7 HARDWARE MAINTENANCE (TR-1) 68
4.7.1 On-site Parts Cache (TR-1) 68
4.8 SOFTWARE SUPPORT (TR-1) 69
4.9 MEAN TIME BETWEEN FAILURE (MTBF) CALCULATION (TR-1) 69

5 FACILITIES INFORMATION 70

5.1 LLNL FACILITIES INFORMATION 70
5.2 LANL FACILITIES INFORMATION 70
5.3 SNL FACILITIES INFORMATION 71
5.4 POWER REQUIREMENTS (TR-1) 71
5.5 COOLING REQUIREMENTS (TR-1) 72
5.6 FLOOR SPACE REQUIREMENTS (TR-1) 72
5.7 DELIVERY REQUIREMENTS (TR-1) 72
5.7.1 SU Installation Time (TR-1) 72

6 PROJECT MANAGEMENT 73

6.1 OPEN SOURCE DEVELOPMENT PARTNERSHIP (TR-2) 73
6.2 PROJECT MANAGER (TR-1) 73
6.3 PROJECT MILESTONES (TR-3) 73
6.3.1 Detailed Project Plan (TR-1) *74*
6.3.2 Tri-Laboratory BASS Final Checkout (TR-1, September 15, 2007) *75*
6.3.3 TLCC07 Phase 1 Build (TR-1, October 2007) *75*
6.3.4 TLCC07 SU Phase 1 Delivery and Acceptance (TR-1, October 2007) *76*
6.3.5 TLCC07 Phase 1 Cluster Integration (TR-1, November 2007) *76*
6.3.6 TLCC07 Phase 1 Option Build (TR-1, October 2007) *77*
6.3.7 TLCC07 SU Phase 1 Option Delivery and Acceptance (TR-1, October 2007) *77*
6.3.8 TLCC07 Phase 1 Option Cluster Integration (TR-1, November 2007) *78*
6.3.9 TLCC07 Phase 2 SU Build (TR-1, January 2008) *79*
6.3.10 TLCC07 Phase 2 SU Delivery and Acceptance (TR-1, January 2008) *79*
6.3.11 TLCC07 Phase 2 Cluster Integration (TR-1, February 2008) *80*
6.3.12 TLCC07 Phase 3 Build (TR-1, April 2008) *80*
6.3.13 TLCC07 SU Phase 3 Delivery and Acceptance (TR-1, April 2008) *81*
6.3.14 TLCC07 Phase 3 Cluster Integration (TR-1, May 2008) *81*
6.3.15 TLCC07 SU Phase 3 Option Delivery and Acceptance (TR-1, October 2007) *82*
6.3.16 TLCC07 Phase 3 Option Cluster Integration (TR-1, November 2007) *83*
6.3.17 TLCC07 Phase 4 Build (TR-1, July 2008) *83*
6.3.18 TLCC07 SU Phase 4 Delivery and Acceptance (TR-1, July 2008) *84*
6.3.19 TLCC07 Phase 4 Cluster Integration (TR-1, August 2008) *84*

7 GLOSSARY 86

7.1 GENERAL 86
7.2 HARDWARE 86
7.3 SOFTWARE 89

1 Background

1.1 Advanced Simulation and Computing Program

The Advanced Simulation and Computing (ASC) Program (formerly know as Accelerated Strategic Computing Initiative, ASCI) has led the world in capability computing for the last ten years. Capability computing is defined as a world-class platform (in the Top10 of the Top500.org list) with scientific simulations running at scale on the platform. Example systems are ASCI Red, Blue-Pacific, Blue-Mountain, White, Q, RedStorm, and Purple (see www.llnl.gov/asc/computing_resources/index_comp_resources.html). ASC applications have scaled to multiple thousands of CPUs and accomplished a long list of mission milestones on these ASC capability platforms. However, the computing demands of the ASC and Stockpile Stewardship programs also include a vast number of smaller scale runs for day-to-day simulations. Indeed, every "hero" capability run requires many hundreds to thousands of much smaller runs in preparation and post processing activities. In addition, there are many aspects of the Stockpile Stewardship Program (SSP) that can be directly accomplished with these so-called "capacity" calculations. The need for capacity is now so great within the program that it is increasingly difficult to allocate the computer resources required by the larger capability runs. To rectify the current "capacity" computing resource shortfall, the ASC program has allocated a large portion of the overall ASC platforms budget to "capacity" systems. In addition, within the next five to ten years the Life Extension Programs (LEPs) for major nuclear weapons systems must be accomplished. These LEPs and other SSP programmatic elements will further drive the need for capacity calculations and hence "capacity" systems as well as future ASC capability calculations on "capability" systems.

To respond to this new workload analysis, the ASC program will be making a large sustained strategic investment in these capacity systems over the next ten years, starting with the United States Government Fiscal Year 2007 (GFY07). However, given the growing need for "capability" systems as well, the budget demands are extreme and new, more cost effective ways of fielding these systems must be developed. This Tri-Laboratory Linux Capacity Cluster (TLCC) procurement represents the ASC first investment vehicle in these capacity systems. It also represents a new strategy for quickly building, fielding and integrating many Linux clusters of various sizes into classified and unclassified production service through a concept of Scalable Units (SU). The programmatic objective is to dramatically reduce the overall Total Cost of Ownership (TCO) of these "capacity" systems relative to the best practices in Linux Cluster deployments today. This objective only makes sense in the context of these systems quickly becoming very robust and useful production clusters under the crushing load that will be inflicted on them by the ASC and SSP scientific simulation capacity workload.

1.2 ASC Capacity Systems Strategy

The ASC "capacity" systems strategy leverages the extensive experience fielding world class Linux clusters within the Tri-Laboratory community. This strategy is based on the observation that the COTs marketplace is demand driven by volume purchases. As such, the TLCC procurement is designed to maximize the purchasing power of the ASC program by changing the past Linux cluster purchasing practices. If we extended existing practices then each of the Tri-Laboratories would separately procure Linux clusters (possibly multiple times

per year) over multiple years. In addition, ASC applications developers and end-users have requested a reduction in the number of different combinations of instruction set architectures, interconnects, compilers, and OS's that they have to support. Thus the Tri-Laboratory community developed a new procurement model based on a common hardware environment over multiple years. However, to balance the risk associated with long term commitment to a single solution in the fast paced commodity space, we have chosen to limit the scope of the procurement to two government fiscal years.

By deploying a common hardware environment multiple times at all three sites over two government fiscal years, it is anticipated that the time and cost associated with any one cluster will be greatly reduced. In addition, it is anticipated that purchasing a huge set of common hardware components will led to lower cost through volume purchases. The Tri-Laboratory site splittings for this procurement in government fiscal year FY07 (4QCY06-3QCY07) and FY08 (4QCY07-3QCY08) have yet to be determined.

The ASC program has also launched a separate parallel activity to determine a Tri-Laboratory productivity on demand user software environment (Tripod). The ASC capacity systems strategy is to minimize the TCO of capacity systems through a common Tri-Laboratory hardware and software environment, in the fullness of time.

1.3 Tri-Laboratory Simulation Environments

The ASC capacity systems will be deployed at the LLNL, LANL and SNL (Tri-Laboratory) sites in the context of existing simulation environments. These simulation environments were developed as part of the ASC program (previously ASCI). Fundamentally, these simulation environments are built around a site-wide global file system (SWGFS) that is shared amongst all of the computing, visualization, networking and storage resources that comprise the simulation environment. LLNL and SNL have standardized their simulation environments on the Lustre parallel file system (www.lustre.org) and LANL has standardized its simulation environment on the Panasas parallel file system (www.panasas.com). Each Laboratory now has a significant investment (>$100M) in their simulation respective environments.

1.3.1 The LLNL Simulation Environment

The LLNL Open Computing Facility (OCF) simulation environment is depicted in Figure 1.3-1-1. A similar simulation environment exists for classified computing (Secure Computing Facility or SCF). The elements of this simulation environment are on the floor at LLNL today. The simulation environment comprises five basic components: 1) 10-40 teraFLOP/s scale Linux clusters; 2) Federated 1 and 10 Gb/s Ethernet networking infrastructure; 3) scalable visualization resources; 4) HPSS based archival resources and 5) Lustre multi-cluster file system components. The Luster file system has three components: client, metadata servers (MDS) and object storage targets (OST). The Lustre client code runs on the compute and rendering nodes of the clusters. The Lustre MDS and OST components are comprised of commodity building blocks and RAID disk devices.

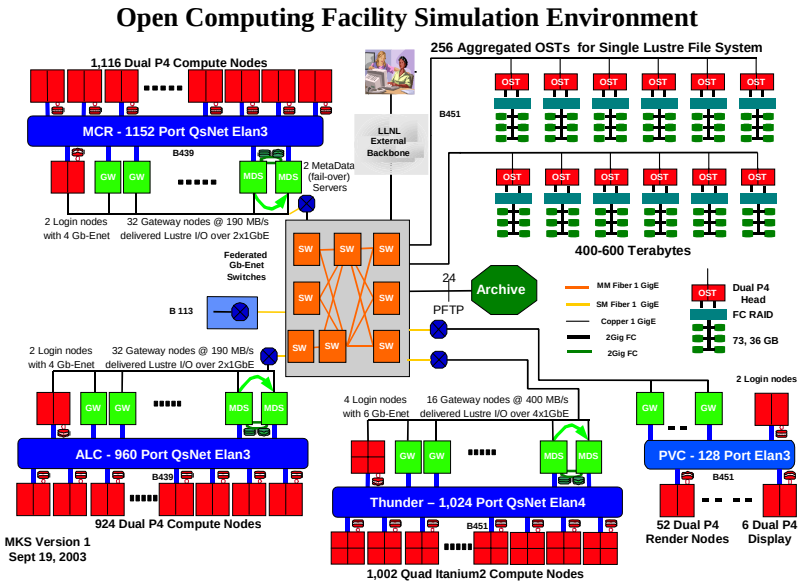

Figure 1.3-1-1: The OCF simulation environment includes multiple 10-40 teraFLOP/s scale Linux clusters, visualization and archival resources. The unifying elements are a federated 1 and 10 Gb/s Ethernet switching infrastructure and the Lustre file system.

1.3.2 The LANL Simulation Environment

The bulk of the LANL simulation capability resides in the secure environment. The Linux portion of the LANL Secure Integrated Computing Network (ICN) is depicted below in figure 1.3-2. Two more similar simulation environments exists for open computing the Open ICN for unclassified open export controlled computing and the Open Collaborative Network (OCN) for unclassified open collaborative computing. The Linux portion of the Secure ICN is in production today running important Weapons Program significant simulations day in and day out. The simulation environment comprises five basic components: 1) 10-70 teraFLOP/s scale Linux clusters; 2) PaScalBB 500 GByte/sec Ethernet networking infrastructure; 3) scalable visualization resources; 4) HPSS based archival resources and 5) Panasas OBSD global parallel file system. The Panasas file system has three components: client, metadata servers and ANSI T10/1355-D object devices. The Panasas client code runs on the compute and rendering nodes of the clusters. The Panasas metadata servers and OBSD's components are comprised of commodity processors, memory, and SATA disks. The Panasas file system implements innovative per object RAID to enable scalable rebuild technology.

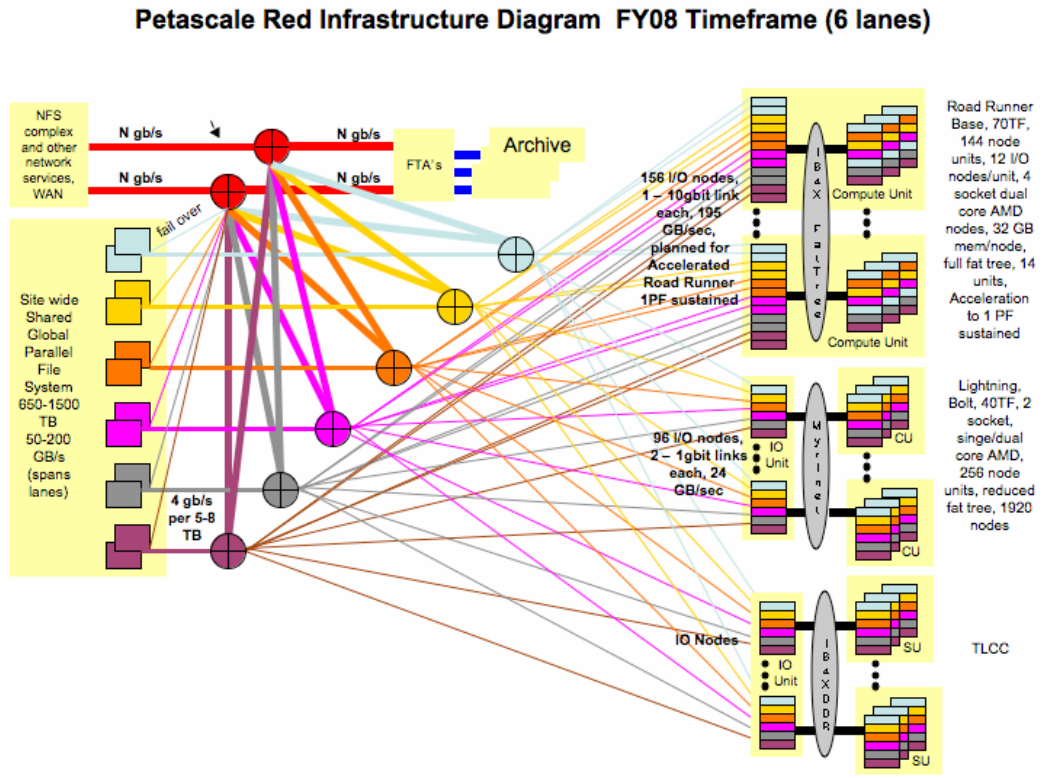

Figure 1.3-2: The LANL Secure ICN simulation environment includes multiple 10 plus teraFLOP/s scale Linux clusters, visualization and archival resources. The unifying elements are an innovative multipath resilient Parallel Scalable Ethernet Back Bone (PaScalBB) and the common globlal parallel Panasas file system (PanFS) based on the ANSI T10/1335-D Object Based Storage Device standard.

1.3.3 The SNL Simulation Environment

The SNL Restricted Network (SRN) simulation environment is depicted in the Figure 1.3-1-2 below. A similar simulation environment exits for classified computing (Sandia Classified Network or SCN). The elements of this simulation environment are on the floor at the SNL California and New Mexico sites today. The simulation environment is comprised of five basic components: 1) 3-60 teraFLOP/s scale Linux clusters; 2) 1 and 10 Gb/s Ethernet networking infrastructure; 3) scalable visualization resources; 4) HPSS based archival resources and 5) Lustre multi-cluster file system components. The Lustre file system has three components: client, metadata servers (MDS) and object storage targets (OST). The Lustre client code runs on the compute and rendering nodes of the clusters. The Lustre MDS and OST components are comprised of commodity building blocks and RAID disk devices. The figure below shows our Lustre environment connected to clusters, archive and visualization clusters.

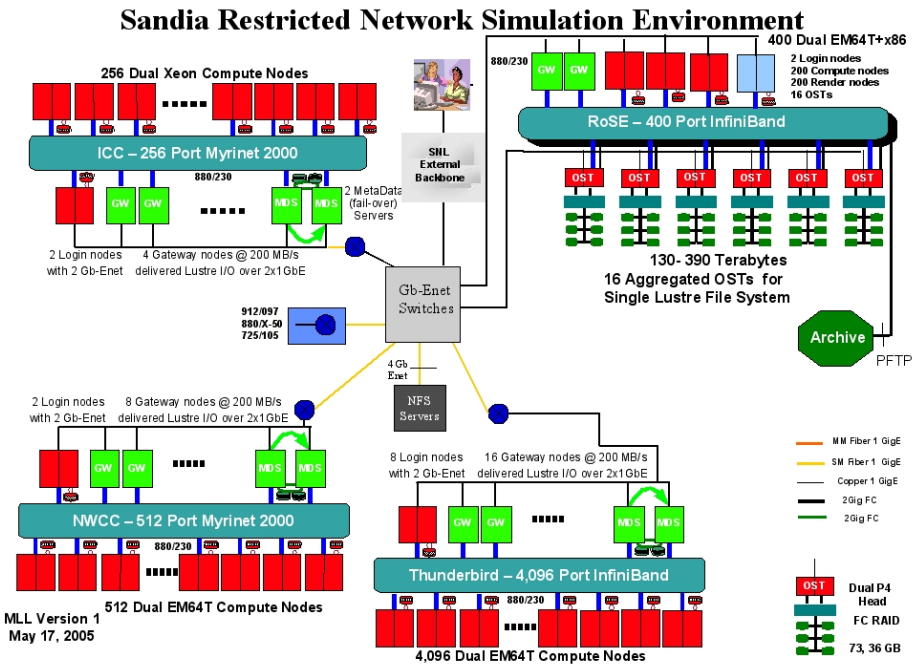

Figure 1.3-1-2: The SNL Simulation Environment is based on Lustre and Linux clusters.

1.4 Utilization of Existing Facilities

The scalable units delivered under this contract shall be delivered to the Tri-Laboratory Sites: The Terascale Simulation Facility (TSF) B453 at LLNL; The SCC facility Building 2327 at LANL and 880/X-50 and 725/105 at SNL Albuquerque and 912/097 at SNL Livermore. The following Sections provide high-level background information to RFP respondents on these facilities.

1.4.1 The LLNL Existing Facilities

An existing facility, portions of the West and East computer floors in B453, will be used for siting the TLCC SU aggregations. See Figure 1.4-1-3. This B453 facility has approximately 48,000ft^2 and >12 MW (7.5 MW for the West floor and currently >4.5 MW for the East floor) of power for computing systems and peripherals and associated cooling available for this purpose. Purple and other systems on the West floor leave approximately 2.5MW available for TLCC SU. In addition, there is about >2.5MW available for TLCC SU on the East floor. SU aggregations can be placed on either floor, but can not span between floors. Facilities modifications to provide the necessary power and cooling for TLCC will need to be accomplished prior to SU delivery. It is therefore essential that the Offeror make available to the University detailed and **accurate** (not grossly conservative overestimates) site requirements for the TLCC SU at proposal submission time. The University will be responsible for supplying the external elements of the power, cooling, and cable management systems.

Figure 1.4-1-3: Building 453 West (left) and West (right) computer floors are being reserved to site TLCC SU aggregations. In this figure north is up.

The TLCC SU aggregations will be migrated to classified operation with access to the SCF networking faculties. In addition, TLCC will be physically located inside a Limited Access Area in a Vault Type Room (VTR). The University will only provide access to the room to authorized personnel under Authorized Escort. All on-site personnel will be required to submit applications for access and be approved by standard University procedures prior to entry into this facility. All on-site personnel will require being DOE P-cleared or P-clearable. RFP responses should indicate if the on-site team has members that are other than U.S. citizens. Physical access to this facility by foreign nationals from sensitive countries (www.llnl.gov/expcon/sensitive.html) will not be allowed. Dialup capability and internet access to the system will be allowed up through acceptance. Authorized individuals may be allowed remote access for running diagnostics and problem resolution. Interaction of the on-site engineering staff with factory support personnel may be limited in some ways (e.g., dissemination of memory dumps from the system may be restricted). These

limitations emphasize the importance of local access to source code, particularly for operating system daemons.

On-site space will be provided for personnel and equipment storage.

They will be expected to practice safe work habits, especially in the areas of electrical, and mechanical.

1.4.2 The LANL Existing Facilities

Existing facilities, portions of Strategic Computing Complex (SCC) and LDCC computer floors, will be used for siting the TLCC SU aggregations. See Figure 1.4-1-32. The SCC will house the bulk of the SU aggregations in the LANL secure environment. The SCC has approximately 303,000ft^2 gross and 44,000ft^2 machine room space with 7.1 MW expandable to 30 MW power and 130,000 GPD expandable to 215,000 GPD water available for secure computational equipment. The LDCC has approximately 12,000ft^2 machine room space for open computational equipment. Facilities modifications to provide the necessary power and cooling for TLCC will need to be accomplished prior to SU delivery. It is therefore essential that the Offeror make available to the University detailed and **accurate** (not grossly conservative overestimates) site requirements for the TLCC SU at proposal submission time. The University will be responsible for supplying the external elements of the power, cooling, and cable management systems. . Before those facility modifications, the available resources are:

	LDCC (rm. 341)	SCC (main computer floor)
Average airflow	600-800 CFM/floor tile	800-1200 CFM/floor tile
Electrical power distribution	208V/120V 3 phase PDUs	208V/120V 3 phase PDUs
Floor loading	250 lb/sqft	300 lb/sqft
Ceiling Height	10 feet	16 feet
Floor depth	24 inch	42 inch

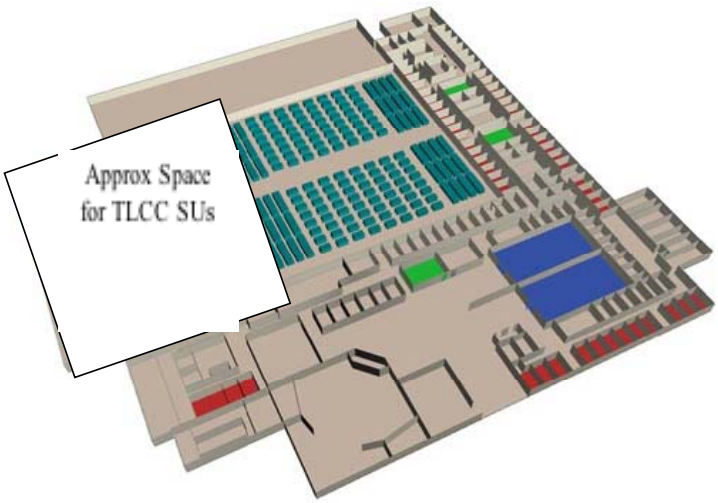

Figure 1.4-2: SCC Building 2327 diagram

Some TLCC SU aggregations will be placed into classified operation with access to the Secure ICN faculties. In addition, TLCC will be physically located inside a Limited Access Area in a Vault Type Room (VTR). The University will only provide access to the room to authorized personnel under Authorized Escort. All on-site personnel will be required to submit applications for access and be approved by standard University procedures prior to entry into this facility. All on-site personnel must be **U.S. citizens**. RFP responses should indicate if the on-site team has members that are other than U.S. citizens. Physical access to this facility by foreign nationals will not be allowed. Dialup capability and Internet access to the system will be not be allowed. Interaction of the on-site engineering staff with factory support personnel may be limited in some ways (e.g., dissemination of memory dumps from the system may be restricted). These limitations emphasize the importance of local access to source code, particularly for operating system components.

On-site space will be provided for personnel and equipment storage.

A safety plan will be required for on-site personnel. They will be expected to practice safe work habits, especially in the areas of electrical, mechanical, and laser activities.

1.4.3 The Sandia Existing Facilities

The existing Sandia facilities, 912/097 at SNL California, and 880/X-50, 880/220, and 725/105 at Sandia New Mexico, can be used for sitting the TLCC SU aggregations. See Figures 1.4-4 through 1.4-7. The SNL California building 912 facility has approximately 1,750 ft^2 (1,250+500 ft^2) for sitting the TLCC clusters. The 912/097 facility has 800 kW of power and 250 tons of cooling available for the TLCC clusters. The Sandia New Mexico 880/X-50 facility has approximately 1,900 ft^2 of floor space, 1,050 kW of power, and 360 tons of cooling. The Sandia New Mexico 880/220 facility has approximately 800 ft^2, 400 kW of power, and 130 tons of cooling. The Sandia New Mexico 725/105 facility has 4,000 ft^2 of floor space, 1,400 kW of power, and 400 tons of cooling capacity for sitting TLCC clusters. SU aggregation can not span floors. Facilities modifications to provide the necessary power and cooling for TLCC will need to be accomplished prior to SU delivery. It is therefore essential that the Offeror make requirements for the TLCC SU at proposal submission time. The SNL will be responsible for supplying the external elements of the power and cooling.

The Figures of the SNL facilities are below. In these figures north is up.

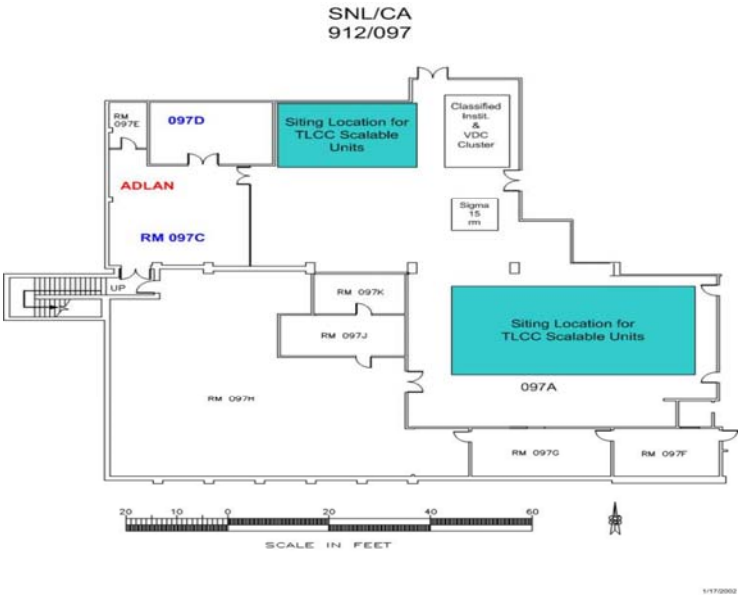

Figure 1.4-1-4: Sandia California Building 912 Computer Room 097.

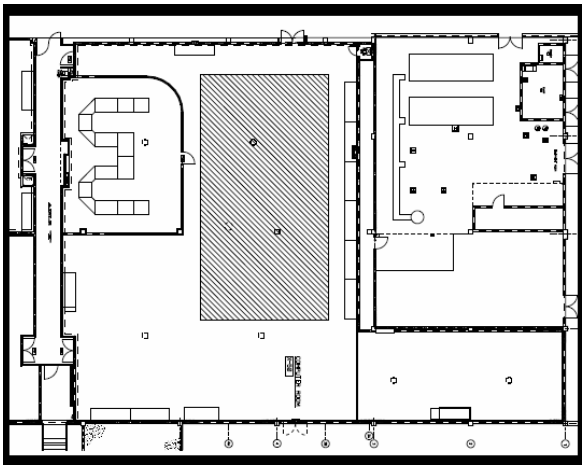

Figure 1.4-1-5: Sandia New Mexico Building 880 Computer Room X-50.

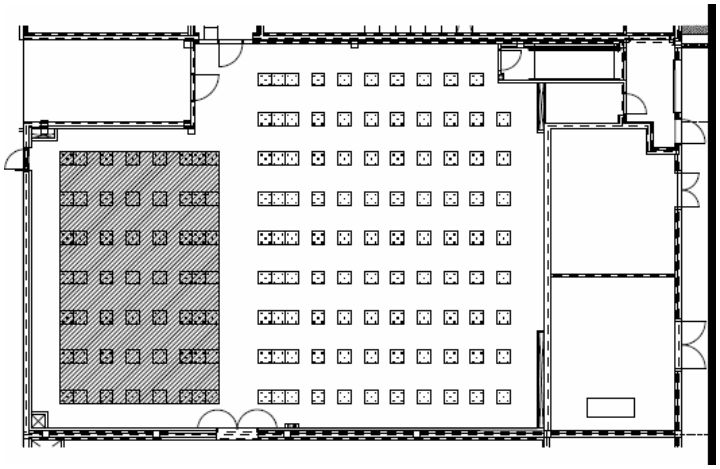

Figure 1.4-6:Sandia New Mexico Building 880 Computer Room 220.

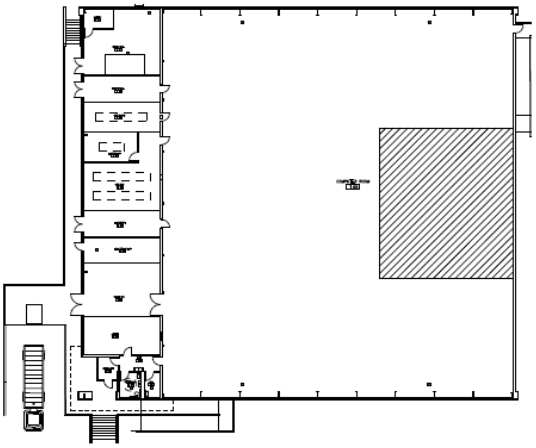

Figure 1.4-7: Sandia New Mexico Building 725 Computer Room.

The TLCC SU aggregations will be migrated to classified operation with access to the SCN networking faculties. In addition, TLCC will be physically located inside a Limited Access Area in a Vault Type Room (VTR). Sandia will only provide access to the room to authorized personnel under Authorized Escort. All on-site personnel will be required to submit applications for access and be approved by standard Sandia procedures prior to entry into this facility. RFP responses should indicate if the on-site team has members that are other than U.S. citizens. Physical access to this facility by foreign nationals from sensitive countries (www.llnl.gov/expcon/sensitive.html) will not be allowed. Internet access to the system will be allowed. Authorized individuals may be allowed remote access for running diagnostics and problem resolution. Interaction of the on-site engineering staff with factory support personnel may be limited in some ways (e.g., dissemination of memory dumps from the system may be restricted). These limitations emphasize the importance of local access to source code, particularly for operating system daemons. On-site space will be provided for personnel and equipment storage. A safety plan will be

required for on-site personnel. They will be expected to practice safe work habits, especially in the areas of electrical, mechanical, and laser activities.

End of Section 1

2 TLCC Scalable Unit Strategy and Architecture

This section describes the overall TLCC Scalable Unit (SU) strategy and architecture.

2.1 TLCC Strategy

As described above, the Tri-Laboratory ASC community requires a large amount of capacity computing resources over the next two fiscal years. In order to affordably and efficiently provide this Production quality computing capacity, the TLCC technical committee has embarked on an approach that extends existing practices within the Tri-Laboratory community while significantly improving TCO. During market survey discussions with the industry, we identified the approach of highly replicated "scalable units" that can be easily built, shipped, sited at the receiving Laboratory and accepted quickly. In addition, there is a balance between the number of SU's deployed and amount of work to maintain a large number of separate clusters at each site. Thus, we have the need to aggregate SU's into clusters built with 1, 2, 4 and 8 SUs. The strategy is to purchase under this RFP all of the components to build the clusters built from SUs and second stage IBA 4x DDR switches.

The delivered SUs in some set of aggregations called clusters will be integrated into existing classified simulation environments at the receiving Laboratory. As such, they need to integrate into the receiving Laboratory's existing multi-cluster file system. For LLNL and SNL, that multi-cluster parallel file system is Lustre from Cluster File Systems, Inc. (www.clusterfilesystems.com) and for LANL that multi-cluster parallel file system is PanFS from Panasas (www.panasas.com/panfs.html).

By replicating the SU many times during the contract we intend to reduce the cost to produce, deliver, install and accept each SU. In addition, this approach will produce a common Linux cluster hardware environment for the Tri-Laboratory user and system administration community and thus reduce the cost of supporting Linux clusters and programmatically required applications on those clusters. However, we require the SU design to be flexible enough to accommodate the following technology improvements:

1. Processor frequency improvements within the same cost and power envelops
2. New processor socket and/or chipset improvements
3. New processor cores
4. SATA disks with higher capacity
5. New memory speed and capacity improvements
6. Interconnect bandwidth and latency improvements

Due to the extremely attractive cost/performance of x86 based Linux clusters and large number of x86 based clusters fielded at all Tri-Laboratory sites and the need for at least 2 GB of memory per processor core, this procurement focuses on solutions that are binary compatible with AMD x86-64 and Intel EM64T with InfiniBand 4x DDR interconnect. However, to meet our improved TCO goals, we have shifted our focus from two-socket nodes found in most past Linux cluster installations, to four socket nodes. Four socket nodes with quad core processor implementations are widely available. This will allow much higher performance Scalable Units for a fixed switch port count. The resulting B:F ratio, measuring the interconnect bandwidth off node (B) to the peak 64b floating point arithmetic performance of the node (F), is well balanced for our applications with this technology choice.

In order to minimize TLCC cluster support costs and the time to migrate a TLCC cluster into classified Production status, the Tri-Laboratory community will supply the Linux cluster software for building, burn-in and accepting the SU's. A Digital Versatile Disk (DVD) will be provided containing the TLCC Build and Accept Software Stack (BASS). BASS consists of a RedHat Enterprise Linux distribution that has been enhanced to support vendor supplied hardware, cluster system management tools required to install, manage and monitor the SU, and a Tri-Lab workload test suite. It is the intent of the Tri-Laboratory community to use the OpenFabrics Enterprise Distribution (OFED) InfiniBand software stack for use on Production computing clusters. More specific details of BASS are provided in Section 2.3 below. The Tri-Lab workload test suite will be used as the SU burn-in, pre-ship test and then run again as an post-ship test after the SU is delivered and assembled at the receiving Laboratory. Once the SU is accepted, the Offeror and receiving Laboratory will be responsible for combining multiple SUs, as directed by the ASC program user community, into clusters with the Offeror supplied spine switches and cables. Final acceptance of these clusters will be accomplished with a scaled up version of the pre/post-ship test.

The support model for TLCC is extremely simple. The Tri-Laboratory community will supply Level 1 and Level 2 support with the Offeror proving Level 3 support functions. On the software side, the Offeror is required to provide and support (level 3) device drivers and other low level specialized software for the provided hardware (including IBA). This support should be provided during normal working hours, Monday through Friday. On the hardware side, the Offeror is required to provide an on-site parts cache of Field Replaceable Units (FRUs) sufficient to cover a at least one week worth of failures without refresh and an Return Merchandise Authorization (RMA) mechanism for return of failed FRUs and refreshing the on-site parts cache. The Tri-Laboratory personnel at each site will provide Level 1 and Level 2 software and hardware support functions that includes responding to problem reports, root cause analysis, reading diagnostics and swapping FRUs. The Offeror shall provide training, on an as needed basis, for hardware FRU replacement over the lifetime of the support period of the contract. Offeror is required to provide each SU with a three year hardware and software maintenance period starting with SU acceptance.

2.2 TLCC Cluster Architecture

This section's description is illustrated by a specific, generic and vendor neutral SU point design. This point design is based on "generic" 2U white boxes with 288 port IBA 4x DDR switches. However, this choice for pedagogical purposed does not constitute a preference by the TLCC technical committee for this solution. The TLCC preference is for an optimized SU design that is more dense (including blades) than this example, yet still meets the facilities limitations for power, cooling and weight. That is, highly optimized and dense solutions (e.g., blades) that are architected in SU in space wasteful way are not seen as beneficial. For clusters larger than 2 SU's, the Offeror may choose to use small port count leaf switches (e.g. 24-port/single switch ASIC switches) to reduce SU costs, to ease node-to-switch integration, and to reduce the maximum number of hops within the IBA fabric.

The ASC Tri-Laboratory's scalable systems strategy for TLCC is exactly the same as that implemented for MCR, Thunder, and Peloton at LLNL (www.llnl.gov/linux/mcr/background/), Lightning and Bolt (www.lanl.gov/news/index.php/fuseaction/home.story/story_id/1473) at LANL and Thunderbird (www.sandia.gov/news/resources/releases/2006/thunderbird.html) at SNL. The

basic idea is that large clusters are usually built from two stage fat-tree interconnects with large radix switches. To architect a production quality cluster in a scalable fashion the system resource like compute nodes, login nodes, management infrastructure and IO infrastructure are divided into smaller groupings associated with each first stage switch. The exact number of each component depends on the size of the switch and the capacity and bandwidth requirements for various components. Thus, a large cluster can be scaled up in these replicated "Scalable Units" (SU). When contemplating purchasing a large number of clusters of various sizes over a two-year period, this approach allows structuring the acquisition and integration activity into a very large number of replicated SUs. This provides the Offeror numerous opportunities to optimize and parallelize SU component purchases, building, testing, shipping, installation and accepting activities.

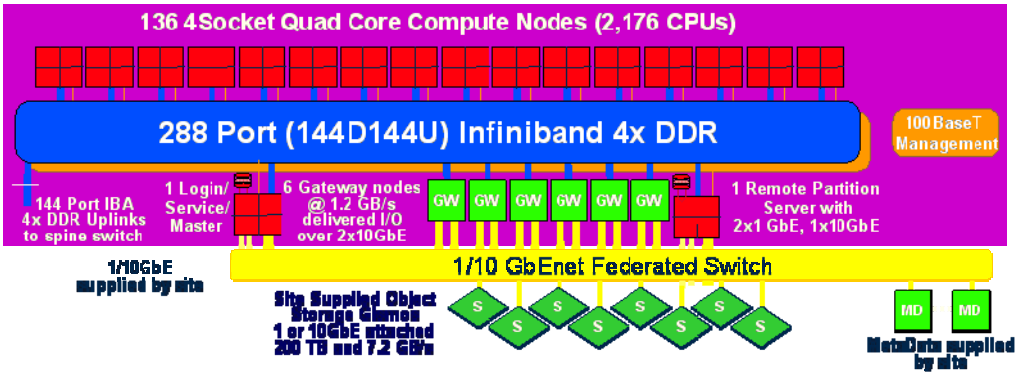

Figure 2-1: TLCC Scalable Unit architecture. Large clusters can be built up by aggregating various numbers of these SU. TLCC cluster architecture includes clustered I/O model, no local node disks, dedicated login/service/master nodes, dedicated gateway nodes and compute nodes all connected to site supplied networking and attached RAID disk resources for Lustre or PanFS.

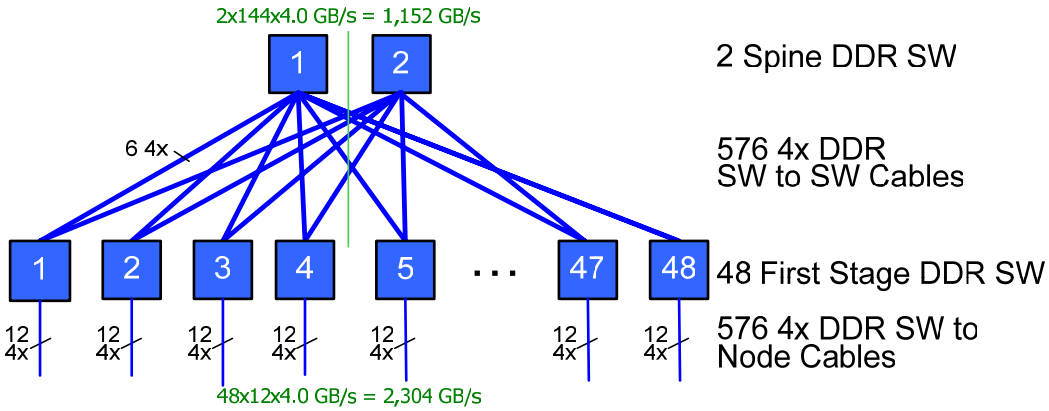

Figure 2-2: InfiniBand 4x DDR interconnect for TLCC SU is based on 24 and 288-port IBA 4x DDR switches. Clusters deployed by Tri-Laboratories will be aggreations of multiple SU configured with a two-stage, full bandwidth, non-blocking, fat-tree federated IBA switch. This 4xSU exmple is based on 48 IBA 4x DDR 24-port first-stage switches configured with 576 4x DDR ports for nodes and 576 4x DDR ports for the second-level

switches (144U144D). The 2 spine switches are each configured with 288 ports connecting the first-level switches (0U288D).

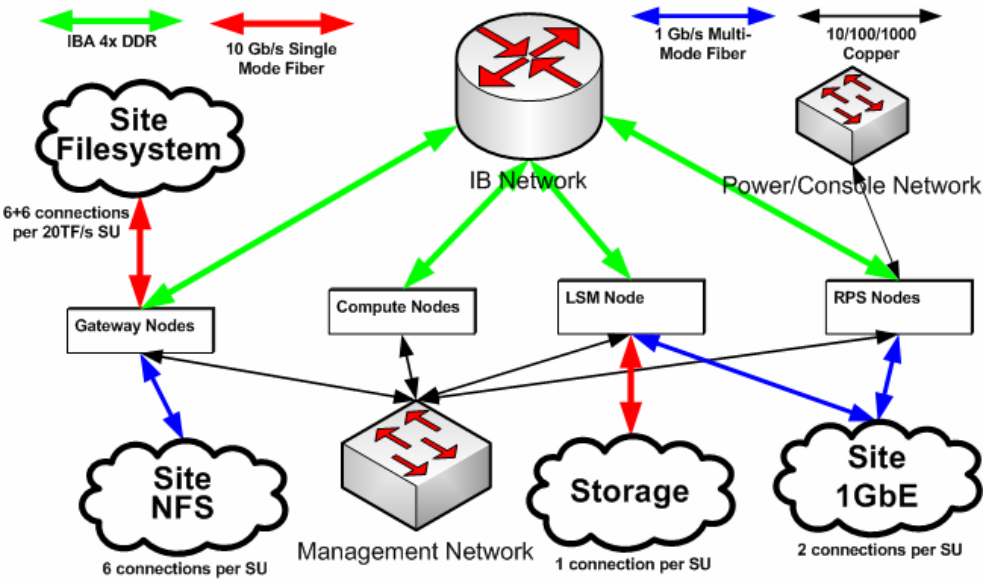

Figure 2-3: TLCC07 Network Layout

The following 20 TF/s 4.6 TB memory SU example in Figure 2-4 is based on 144 2U/4-Socket nodes with IBA single port 4x DDR PCIe 8x card. The SU has 7.2 GB/s global I/O bandwidth through the six gateway nodes. This SU requires 7x42U compute racks for compute nodes and terminal server and 2x1U 24-port IBA switches and 1x42U IO rack for a combination of compute nodes, gateways, login/service/master and remote partition server nodes, terminal server and the management Ethernet switch.

The 8x42U rack SU example in Figure 2-4 is about 16' long and 3' wide. Air flows from the front of the row to the back. Within the SU, the distances between all the nodes and the IBA switch are much less than 10m and so 144 CX4 copper cables are used between the nodes and the IBA switch. The compute nodes each require about 700 watts and hence each compute rack is about 12.6 KW and weighs about 800 lbs or about 133 lbs/ft^2.

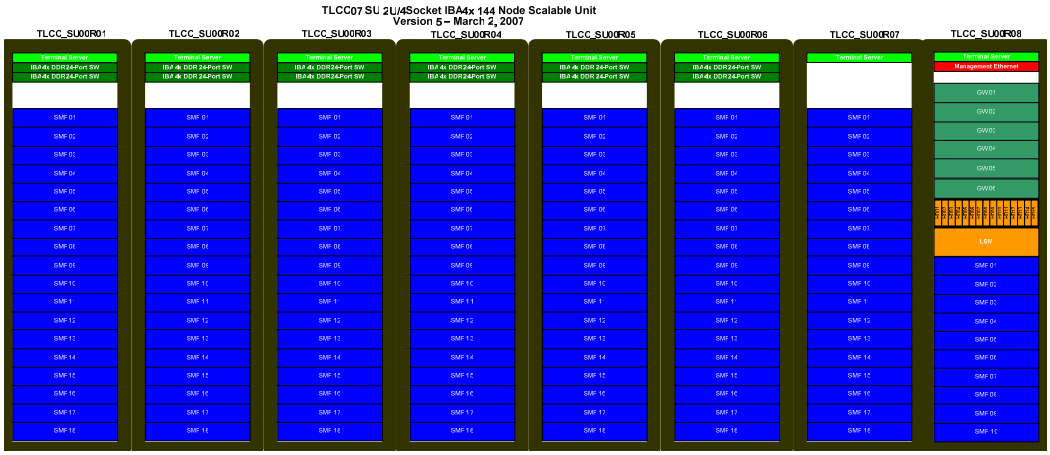

Figure 2-4: TLCC SU eight rack layout based on 144 2U/4Socket nodes and 12 24-port IBA 4x DDR switches.

After the SU are delivered, installed and accepted, the receiving Laboratory and the Offeror will combine multiple SUs together to form a TLCC Cluster. The floor plan layout in Figure 2-5 shows a hypothetical installation of four clusters of four SU each (a total of 16 SU). Each cluster of four SU is in a separate row with the two 288 port IBA 4x DDR spine switches in Red rack the middle of the row. The two SU on the left of the spine switch rack have the rack layout shown in Figure 2-4. The two SU on the right of the spine switch have the rack layout reversed (mirror image). This is done in order to minimize the distance between the first stage switch and the second stage switch.

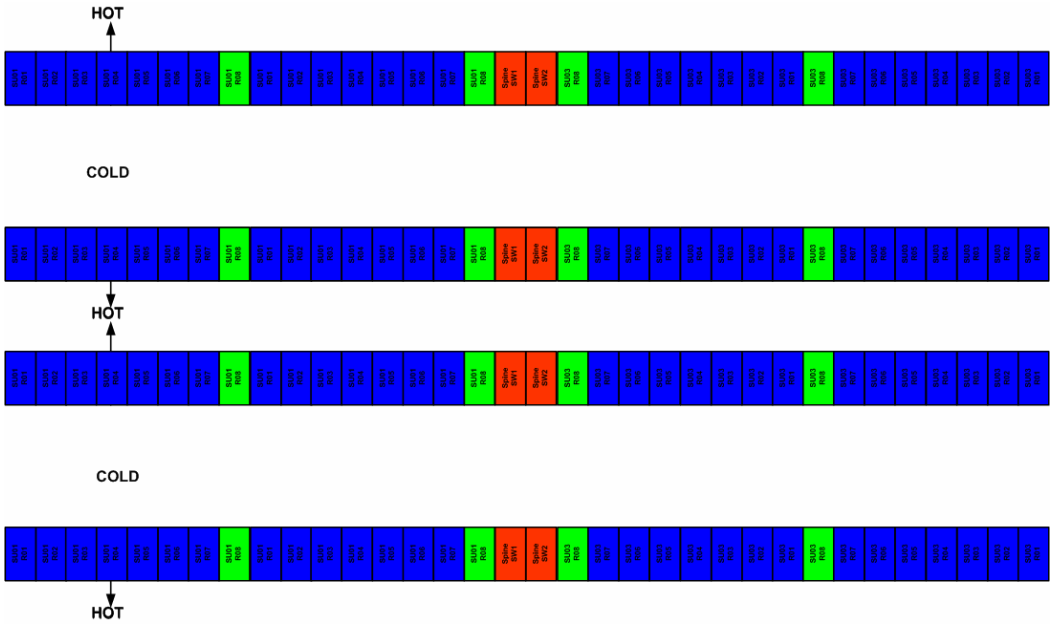

Figure 2-5: Hypothetical TLCC 4x4SU Cluster Layout. Each cluster is based on 4 SU with 8 racks per SU and 1 rack of 2 IBA 4x DDR Spine Switches. Note that the 144 IBA 4x

DDR links between L1 Switch and Spine Switch are Copper. There are no IBA links between rows. Each Cluster has only 1 and 10Gb/s Ethernet external links

With this layout, cable distances between the SU IBA switch and the IBA spine switches are less than 10m and so it is possible to use copper cables. No IBA cables are required between rows because each row is an independent cluster. In addition, each cluster has four login/service nodes each with two 1 Gb/s Ethernet and one 10 Gb/s Ethernet connection to the Laboratory Ethernet infrastructure. In addition, each cluster has twenty-four gateway nodes each with two 10 Gb/s Ethernet connection to the Laboratory Ethernet infrastructure.

Note in Figure 2-5 that the cold isles are wider than the hot isles and the airflow through the racks emanates from floor tile grates in front of the racks in the cold isle. Air exhausts from the racks into the hot isle and is removed from the room via grates in the ceilings above the hot isles. There are no air handlers on the floor at LLNL or LANL, but will be present at SNL. In addition, power is provided to the racks by cables running under floor from wall panels. In this example, there are no PDUs on the floors. To minimize facilities modifications costs and breaker utilization, TLCC racks should use a minimum number of circuits (one) and maximize the utilization of that circuit (up to the maximum of 80%).

2.3 TLCC Software Environment

To execute the ASC capacity systems strategy, the TLCC SU need to be transitioned to classified Production service as quickly after acceptance as possible. In order to facilitate this, each receiving Laboratory will install their production Linux distribution on each SU after acceptance. This software will be targeted to the Offeror's hardware environment in collaboration with the Offeror after contract award and prior to the SU manufacture.

For the purposes of the TLCC project, the Tri-Laboratory community will create a **single** Build and Accept Software Stack (BASS). The BASS combines elements of each Laboratory's Linux distribution and is focused on providing the Offeror with the software that will be needed to build, debug and validate each SU, along with the suite of applications that will be required for acceptance. This software will be tailored to the Offeror's hardware environment in collaboration with the Offeror after contract award and prior to the SU manufacture.

Section 2.3.1 below outlines the contents of BASS, followed by a description of the Linux distributions used at LLNL in section 2.3.2; at LANL in section 2.3.3 and at Sandia in section 2.3.4. As previously mentioned, the BASS is a subset of the more full-featured distributions used in production by the Laboratories. It is important for the Offeror to understand the eventual software environments that will be deployed by the Laboratories once the SU are accepted. Over time, those software environments will evolve according to Tripod guidelines (see section 1.2).

2.3.1 Tri-Laboratory Build and Accept Software Stack (BASS)

Distribution – BASS is derived from RedHat Enterprise Linux (RHEL) version 4 and later. The details of the BASS software will be worked out with the Offeror as input to Milestone 6.3.1.

Kernel – BASS replaces the RHEL kernel with the latest stable kernel.org kernel (version 2.6.x). This kernel will include patches to support InfiniBand as well as a number of other kernel modifications deemed necessary by the Tri-Lab community

InfiniBand Stack – The OpenFabrics OpenFabrics Enterprise Distribution (OFED) stack will be provided as part of BASS. The OpenSM InfiniBand subnet manager will be included as part of BASS. The Offeror is expected to use OpenSM as the fabric subnet manager throughout the acceptance test stage. The Tri-Labs will work with Offeror to incorporate this InfiniBand stack into the BASS.

Device drivers – Any device drivers required to support the Offeror's hardware, which are not available in the RHEL standard distribution, should be provided by the Offeror for incorporation into BASS. This additional or modified software must be provided in the form of buildable source RPMs.

Diskless cluster installation and configuration – Offeror Supplied…

Remote Management - If Offeror proposes an IPMI solution, node firmware required to implement the IPMI 2.0 protocol should be provided by the Offeror. FreeIPMI will be used to validate basic IPMI 1.5 and IPMI 2.0 compliance. PowerMan (www.llnl.gov/linux/powerman/powerman.html) will be provided for remote power management. Powerman will include support for IPMI if the Offeror's solution requires it. ConMan (www.llnl.gov/linux/conman/conman.html) will be provided if Offeror proposes a traditional console management solution. If Offeror proposes IPMI Serial over LAN (SOL), solution is then an interface ConMan can utilize will be provided.

Cluster Monitoring – Ganglia and/or Supermon and/or SNMP will be used to gather data of interest from the nodes (in-band). This data may include node resource utilization such as CPU, memory, I/O, runaway processes, etc. Node environmental data (e.g. temperature, fan speeds, voltages) will be collected by utilizing the LMSENSORS kernel module (with interface to IPMI if appropriate).

Resource Manager – Torque (www.clusterresources.com/products/torque)

Batch Scheduler - Maui (www.clusterresources.com/products/maui)

Vendor Compilers - Must include at least one major Fortran compiler as provided by either Intel, Pathscale, or PGI compilers.

Test Harness – Gazebo (LANL)

BIOS management tools - Any BIOS management tools required to support the Offeror's hardware should be provided by the Offeror for incorporation into BASS

2.3.2 LLNL Clustered High Availability Operating System (CHAOS)

Livermore Computing (LC) produces and supports CHAOS (http://www.llnl.gov/linux/chaos), a cluster operating environment for HPC Linux clusters. At the core of CHAOS 3.0 is the Red Hat Enterprise Linux 4 (RHEL4) distribution. Some

components of the RHEL distribution are modified to meet the demands of high-performance computing installations, operations and support. LLNL has added a number of additional cluster-aware components to RHEL.

A CHAOS distribution contains a set of RPM (Red Hat Package Manager) files, RPM lists for each type of node (compute, management, Gateway, and login), and a methodology for installing and administering clusters. It is produced internally and therefore supports a short list of hardware and software. This approach permits LC to support a large number of clusters with a single CHAOS release supported with a small staff and to be agile in planning its content and direction.

In addition to the products of Open Source development described below and the base Red Hat Linux distribution, CHAOS includes the following software components:
Kernel — The CHAOS kernel is based on a Red Hat kernel with additions in the areas of device support for IBA, VFS modifications for Lustre, ECC and FLASH memory device support for Intel motherboard chipsets, crash dump support, miscellaneous bug fixes, and optimized configurations for LC's hardware.
OFED — The Open Source OpenFabrics Enterprise Distribution (OFED) IBA software environment used to run parallel programs over the InfiniBand fabric.
crash — The MCL crash dump analysis suite is used to examine post mortem contents of a kernel crash dump.
lm_sensors — Hardware monitoring, linked to the SNMP host monitoring system, monitors motherboard chipset sensors such as temperature, fan speed, and power supply voltages.
fping — Fping is a rudimentary node status tool that can ping nodes in parallel. In combination with genders tools, fping can quickly find nodes in the cluster that are turned off or otherwise unreachable on management network.
OpenSSH — OpenSSH provides encrypted remote login/shell service that is PAM-aware.
PAM Tools — PAM modules for One Time Passwords (OTP), Kerberos V, and RMS (to authorize a user's access to a compute node only when RMS is running that user's job) are used to leverage LC's DCE and OTP infrastructure for PAM-aware applications.
Firmware — Firmware images for motherboards, including FLASH/CMOS support software, is included in CHAOS. Firmware and support software for power control/serial console hardware is also included.
MPI Test Suite — The Pallas MPI Benchmark (PMB), Effective Bandwidth test (BEFF), and MPI ping-pong test (mping) are packaged with CHAOS with a script to maintain a continuous MPI workload under RMS for testing purposes.
NTTCP—The NTTCP TCP bandwidth test is packaged along with genders-aware scripts that can simulate load on the management Ethernet for testing purposes.
sudo — Sudo extends administrative privileges to non-root users.
Intel Compilers — The Intel EM64T FORTRAN, C, and C++ compilers.
PGI Compilers — The Portland Group Fortran and C Compilers
TotalView — The TotalView parallel Debugger from Etnus.
Other Libraries/Tools — Other libraries and tools such as Atlas, COG, NDF, netCDF, Hyper, ScaLAPAC, OpenGL, Yorick, Silo, VTK, and Findentry are maintained on LC Linux systems.

2.3.2.1 LLNL Cluster Tools

The following Open Source cluster tools (http://www.llnl.gov/linux/ctp/) are under active development and have been deployed on all LC Linux clusters:

pdsh — The Parallel Distributed SHell utility executes processes across groups of nodes in parallel. It is also capable of running small MPI jobs on the Elan interconnect.

YACI — Yet Another Cluster Installer is Livermore's system installation tool based on various cluster installers such as VA system imager and LUI.YACI can fully install the 1,152-node MCR cluster in about 15 minutes. It is image-based and can use either an NFS pull or multicast mechanism to install many nodes in parallel.

Genders — Genders is a static system configuration database and rdist Distfile preprocessor. Each node has a list of "attributes" that in combination describe the configuration of the node. The genders system enables identical scripts to perform different functions depending on their context. An rdist Distfile preprocessor expands attribute macros into node lists allowing very concice Distfiles to represent many large clusters.

ConMan — The ConMan console manager (http://www.llnl.gov/linux/conman) manages serial consoles connected either to hardwired serial ports or remote terminal servers (telnet based). Performs logging of console output, and manages interactive sessions, permitting console sharing, console stealing, console broadcast, and interfaces for transmitting a serial break or resetting a node via PowerMan. ConMan does not directly support IPMI, but will be modified to interface a daemon which converts IPMI serial-over-LAN to telnet protocol.

PowerMan — The PowerMan power manager (http://www.llnl.gov/linux/powerman) manages system power controllers and is capable of sequenced power on/off for groups of nodes and initiating reset (both plug off/on and hardware reset if available). Powerman currently supports the Linux NetworX ICE box, WTI RPC's, IPMI 1.5 through ipmipower component of FreeIMPI and the API Networks modified Wake-on-lan. It can be extended to support new hardware. PowerMan does not directly support IPMI 2.0, but ipmipower will be modified to utilize IPMI 2.0.

Host Monitoring System — TLCC will monitored in-band (while Linux is running) by polling via NET-SNMP (http://net-snmp.sourceforge.net/). Among the information polled is motherboard sensor information and information about failing hardware devices such as memory and disks. In addition, PowerMan can extract out-of-band monitoring information such as case temperature from some remote power control devices that have this capability. LC's SNMP based host monitoring system stores current state in a MySQL database and long-term state in an RRD (round robin database). Collection software polls cluster nodes in parallel using SNMP bulk queries and a sliding window algorithm to reduce polling latency. Status is presented via web using Apache and PHP.

2.3.2.2 Simple Linux Utility for Resource Management

SLURM is an Open Source, fault-tolerant and highly scalable cluster management and job scheduling system for clusters containing thousands of nodes. SLURM is the production resource manager on all LC Linux clusters that use the Quadrics Elan3 interconnect, and it has been ported to other networks as described below (http://www.llnl.gov/linux/slurm/).

The primary functions of SLURM are:
Monitoring the state of nodes in the cluster.

Logically organizing the nodes into partitions with flexible parameters.
Accepting job requests. While SLURM can support a simple queuing algorithm, Moab Cluster Suite will manage the order of job initiations through its sophisticated algorithms described in Section 2.2.4 of this document.
Allocating both node and interconnect resources to jobs.
Monitoring the state of running jobs, including resource utilization rates.

SLURM utilizes a plug-in authentication mechanism that currently supports authd and the LLNL-developed munge protocol. The design also includes a scalable, general-purpose communications infrastructure. APIs support all functions for ease of integration with external schedulers. SLURM is written in the C language, with a GNU autoconf configuration engine. While initially written for Linux and Quadrics Elan3 interconnects, the modular design allows for ease of portability. A vendor partner collaborator has ported SLURM to IBA and Myrinet. LC has ported SLURM to Quadrics Elan4, IBM BlueGene/L and IBM pSeries Federation and Ethernet based clusters.

2.3.2.3 Moab Cluster Suite

Moab Cluster Suite is a professional cluster management solution that integrates scheduling, managing, monitoring and reporting of cluster workloads. Moab Cluster Suite simplifies and unifies management across one or multiple hardware, operating system, storage, network license and resource manager environments to increase the ROI of cluster investments. Its task-oriented graphical management and flexible policy capabilities provide an intelligent management layer that guarantees service levels, speedy job processing and easily accommodates additional resources. For more information see http://www.clusterresources.com/pages/products/moab-cluster-suite.php

2.3.2.4 Lustre Cluster Wide File System

LC plans to utilize the Lustre Cluster Wide File System on clusters built up from the TLCC SU. To that end, LC and Cluster File Systems, Inc., have been actively engaged in developing a production version of Lustre to run on MCR since the summer of 2002. See http://www.lustre.org/ for more information on Lustre. Currently, Lustre is in production status with MCR, ALC, Thunder, and Lilac. LC is also working with CFS to port Lustre to IBM BlueGene/L.

2.3.2.5 The Livermore Computing Linux Cluster Support Model

LC Open Source developers (Cluster Tools, LCRM, SLURM, Lustre) work closely with system administrators and users to resolve problems on production systems. For any given software package, there is a designated package owner who handles any support issues that arise in production. Depending on the nature of the package, owners may be the primary developer and fix bugs themselves, or they may be the liaison to an external support resource.

External support relationships are primarily developer-to-developer. In the case of Red Hat, we have a full-time Red Hat engineer on site who works directly with Livermore systems and support people and acts as the liaison to Red Hat for everything in the Red Hat Linux distribution.

2.3.2.6 Integration Testing

Each CHAOS release is subject to integration testing that includes regression tests for past problems, basic functionality tests, and real users' applications. Each of the software components is developed asynchronously, but come together in system (CHAOS) releases and separate package (DPCS, Lustre Lite, SLURM, Cluster Tools) releases. Due to this separation, system and package testing and installation on production clusters can be scheduled and executed independently. The LC Linux testbed has a number of small (average 16 nodes) clusters that are available for unit testing of individual software components, integration testing of a complete CHAOS release, and debugging of defects that arise in production. These development clusters, along with the project CVS repository, can accommodate external collaborators working with LC on the Open Source projects described above.

2.3.3 LANL Linux Capacity Environment (LLCE)

LANL is participating in the Tripod project (see section 1.2) to develop a seamless software environment for use by the NNSA Tri-lab community (Los Alamos, Livermore, and Sandia national laboratories), initially targeted at Linux capacity computing clusters. Tripod recommendations will evolve over time. Currently, in addition to the products of Open Source development described below and the base Red Hat Linux distribution, the LANL software stack includes the following software components:

Kernel — The kernel(s) that are utilized within the LANL environment are based directly on kernels from www.kernel.org. We use a standard distribution of the kernel, rather than depending on any one vendor supplied software stack.
OpenFabrics — The Open Source IBA software environment used to run parallel programs over the InfiniBand fabric. Please see www.openfabricshttp://www.openib.org/.org for more information.
LA-MPI (OpenMPI) — LANL is one of the original founding members of OpenMPI. OpenMPI is a combination of technologies available from FT-MPI,LA-MPI,LAM/MPI and PACX-MPI. Please see www.open-mpi.org for more information or http://public.lanl.gov/lampi
MPICH — Many vendors supply MPICH as a part of their distribution. We maintain a copy of MPICH on systems at LANL. Please see www-unix.mcs.anl.gov/mpi/mpich
ROMIO — A high performance, portable MPI-IO implementation, which is highly utilized at LANL. Please see www-unix.mcs.anl.gov/romio
LSF — LANL is starting to replace LSF with Moab and Torque. LSF is software for managing and accelerating batch workload processing for compute-and data-intensive applications. With LSF you can intelligently schedule and guarantee completion of batch workload across a distributed, virtualized environment. LSF and BProc have been integrated together in the LANL environment. Please see www.platform.com/product/LSF
Clustermatic — Clustermatic is a collection of new technologies being developed specifically for new cluster architectures. Each technology can be used separately, and thus does not preclude integration with other clustering efforts or even other types of computing environments.

- **Philosophy**
  - Fewer parts means fewer sources of problems
  - Silicon focus: CPU, RAM, network required, disks not essential
  - Compute nodes do not need full complexity OS

- Quick reboot: Enables guaranteed node cleanup after each job
- Need to know: Nodes see only application relevant portions of the filesystem
- Modular
- Easy to build/replicate
- Supports 2 – 2048 nodes (and counting)
- System administrations is simplified
  - Upgrade OS within entire system with one command
  - No need to synchronize versions
- Entire software suite GPL
- Uses Beoboot to manage booting of entire cluster

- Current elements of Clustermatic are: BProc, LinuxBIOS, SuperMon, Beoboot and Guard. Please see www.clustermatic.org

LinuxBIOS — LinuxBIOS replaces the normal BIOS bootstrap mechanism with a Linux kernel that can be booted from a cold start. Cluster nodes can now be as simple as they need to be, perhaps as simple as a CPU and memory, no disks and no filesystem. This allows for nodes to be booted and up and running in less than 3 seconds. LinuxBIOS is the BIOS of choice in LANL's Linux clusters, lab wide. Please see www.linuxbios.org

BProc — LANL is currently using BProc but TLCC systems will use system software consistent with Tripod recommendations. BProc (Beowulf Distributed Process Space) provides a single system image of the entire cluster. LinuxBIOS nodes come up autonomously and contact the "front end" node which sends them a BProc kernel to boot and registers them as part of the cluster. Users run programs on the front end, which then migrates the jobs to the other cluster nodes. BProc itself consists of a small set of kernel modifications, utilities and libraries which allow a user to start processes on other machines in a cluster (including reboot). Remote processes started with this mechanism appear in the process table of the ftont end. This allows remote process management using the normal UNIX process control facilities. Signals are transparently forwarded to remote processes and exit status is received using the usual wait() mechanisms. BProc allows for a single point of process management. Please see bproc.sourceforge.net

Supermon — Supermon is a high speed cluster monitoring tool that can collect 1000 samples per node per second without noticeable affect on the cluster nodes. The data from Supermon can be used to monitor node health and perform remote node maintenance. In addition, the monitoring information can be used to predictively react to node failures. Please see supermon.sourceforge.net

lm_sensors — Hardware monitoring, linked to the SNMP host monitoring system, monitors motherboard chipset sensors such as temperature, fan speed, and power supply voltages. Please see www2.lm-sensors.nu/~lm78/

V9FS — V9FS is a private namespace filesystem, developed at LANL. It offers:
Avoids problems associated with global mounts
Processes are provided with a private shared filesystem
Namespace exists only for the duration of the process
Nodes are returned to "pristine" state once the process is complete
Please see v9fs.sourceforge.net

XCPU — XCPU is a network filesystem interface based on the Plan 9 9P protocol. It provides:

A unified information transfer mechanism that is not operating system or platform dependent, enabling control of multiple machines and architectures. Supported platforms include Linux, MacOS X and the Cell architecture.
Fast scalable process execution/startup within a private shared filesystem to remote compute nodes from multiple master nodes.
Please see http://plan9.bell-labs.com/sources/contrib/iwp9-2006/papers/xcpu-madrid.pdf
fping — Fping is a rudimentary node status tool that can ping nodes in parallel. In combination with genders tools, fping can quickly find nodes in the cluster that are turned off or otherwise unreachable on management network. Please see www.fping.com
OpenSSH — OpenSSH provides encrypted remote login/shell service that is PAM-aware. Please see www.openssh.org
PAM Tools — PAM modules for One Time Passwords (OTP), Kerberos V. Please see http://www.kernel.org/pub/linux/libs/pam/
Firmware — Firmware images for motherboards, including FLASH/CMOS support software, are maintained at LANL. Firmware and support software for power control/serial console hardware is also maintained.
MPI Test Suite — The Pallas MPI Benchmark (PMB), Effective Bandwidth test (BEFF), and MPI ping-pong test , BW a bandwidth/pattern sending program are packaged and maintained for testing purposes at LANL.
Intel Compilers — The Intel EM64T FORTRAN, C, and C++ compilers.
PGI Compilers — The Portland Group Fortran and C Compilers
PathScale Compilers — Please see www.pathscale.com
TotalView — The TotalView parallel Debugger from Etnus.
Other Libraries/Tools — Other libraries and tools such as Atlas, NDF, netCDF, HDF, Hyper, ScaLAPAC, OpenGL and VTK are maintained on LANL Linux systems.
System Initialization — System Imager which provides a capability to deploy our system software on new systems or rebuild failed systems. (www.systemimager.org)
Software Management — Cfengine which provides/maintains a database of our software and allows for automation (i.e., scripting) of routine system administration tasks. From the www.cfengine.org,

- "Cfengine, or the configuration engine is an autonomous agent and a middle to high level policy language and agent for building expert systems to administrate and configure large computer networks. Cfengine is designed to be a part of a computer immune system. It is ideal for cluster management and has been adopted for use all over the world in small and huge organizations alike." First developed in 1993, cfengine provides us a capability of integrating our customer specialized software needs via RPMs that are not part of a standard Linux distribution.
- Tripwire (www.tripwire.org) – monitor the occurrence of software changes and alerts us to "violations" of our FPiCM philosophy. A side benefit is that it also alerts us to "suspect" operations on our systems.
- rsync (samba.anu.edu.au/rsync)– provides a secure and reliable method for deploying software via the network to production systems
- rpm (www.rpm.org) - provides an easy and well-known method for installing, updating, or removing software. Use of RPMs is consistent with the Linux Standards Base (www.lsb.org) goal
    - "to develop and promote a set of standards that will increase compatibility among Linux distributions and enable software applications to run on any compliant system. In addition, the LSB will

help coordinate efforts to recruit software vendors to port and write products for Linux."

- LSB is rapidly being accepted by commercial software vendors in deploying their software which reduces the amount of effort LANL incurs when deploying updates to commercially provided software

2.3.3.1 Panasas PanFS Multi-Cluster Environment Wide Global Parallel OBSD Based File System

LANL plans to utilize the centralized and globally shared Panasas File System (PanFS) on clusters built from the TLCC SU. See www.panasas.com for more information on PanFS. Currently, LANL has three PanFS global parallel file systems in production, one in each computing environment, secure, open and collaboration networks. PanFS is used in production weapons simulation computing daily and by all major Linux supercomputers at LANL, as well as many commercial production facilities. The following image is an example of how the systems might be connected.

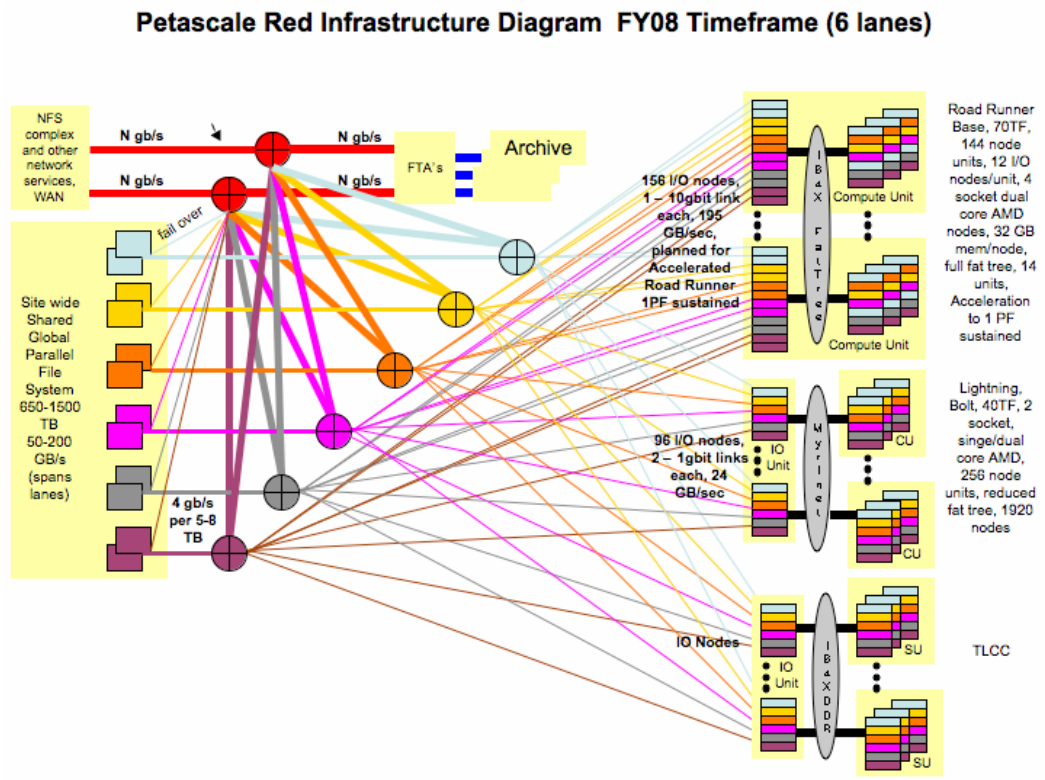

Figure 2.3-2-6: Expected Secure Linux Environment with TLCC.

2.3.3.2 Regression analysis and Testing at LANL

Each tool/package provider does individual testing and regression (like PanFS, Lampi,etc.) We have Storming (and like systems) where new integrated composite images are tested/regression tested as a total system, with our LUP tool which has depth tests for each area, MPI, file systems, ….., and cross cutting broad synthetic tests, and

actual user apps (sage, qcd, sweep). We have a locally developed Gazebo harness for these tools which runs the test/regression and reports and keeps stats/db etc.

2.3.3.3 Configuration Control

It is crucial that there be one true copy of the source code. There are several suppliers that provide a configuration management capability, access to the source repository, as their product or part of it.

Many developers feel it is also necessary to have a single repository to track defects and enhancement requests, documents, mail lists, developer assignments, and more.

CVS (Concurrent Versions System) and SVN (Subversion) are provided on all production systems. Repositories can be created on and accessed from any production system.

SourceForge is currently provided in the Yellow and Red networks. SourceForge adds many collaborative development components, such as trackers, document management, task list management, mail lists, and newsgroups.

2.3.3.4 Parallel Programming Models & Run-Time Systems

The current parallel programming models being used by the ASC program are MPI, threads, OpenMP, and hybrids of these. Interoperable libraries that fill these roles are required.

LA-MPI (or OpenMPI on newer systems) is provided on all systems. If a vendor-supplied MPI implementation exists, it is also provided.

All systems have POSIX compliant thread libraries installed.

OpenMP is provided as an extension to at least one compiler on each system. This coupled with the POSIX compliant thread libraries, and thread-safe MPI implementations allow for hybrid parallel programs to be built and run.

2.3.3.5 Parallel I/O Libraries

The I/O libraries must support all modes of writing files for parallel programs. Specifically, where one process reads/writes many files; many processes read/write one file; and many processes read/write many files.

LANL provides the standard I/O libraries that come with the various compilers. Versions 4 and 5 of HDF are also provided. The applications may use these to build their I/O capabilities.

The MPI implementations provide MPI/IO to allow for serial and parallel I/O.

UDM (Unified Data Model) is provided to address the applications' needs for data organization and parallel I/O.

2.3.3.6 Compilation/Build Environment

It is necessary to provide all the tools to turn source code into executable applications. Certain languages, build tools, and scripting capabilities are required to accomplish this task. LANL provides these GNU utilities: autoconf, automake, diffutils, fileutils, findutils, gawk, gmake, gnu m4, and sed. Parallel make is provided via the *gmake –j* command.

On 64-bit Linux, LANL provides the PGI C/C++/Fortran 77/90, PathScale C/C++/Fortran 77/90/95, Intel C/C++/Fortran 77/90/95, and GCC C/C++/Fortran 77/gfortran. The GCC Fortran 77 compiler does not support Cray pointers, and the GCC gfortran is not yet widely deployed. The GCC compilers do not support OpenMP.

LANL provides the Java language and development kit on all of its systems.

LANL provides common scripting languages and UNIX shells. LANL does not provide ksh or expect.

2.3.3.7 Debugging and Correctness Testing

These are tools to make sure the applications are doing the correct thing when they execute.

TotalView fills most of our needs in graphical, interactive debugging of parallel programs. A concern with Etnus is its vulnerability to being acquired by a larger company, like Intel, and the product becomes available on Intel platforms only. LANL has installed TotalView on its systems. 1/3 of LANL's largest system is 1,365 processors. At this level, it works with some quirks that are being addressed by Etnus. Some areas of development in parallel debugging have been identified, but are not actively being addressed.

There are really a couple of issues here. Memory usage deals with how much of various kinds of memory are being used at points in the program. Memory correctness deals with errors in how memory is allocated and released, or referenced as the program runs. LANL has a few tools available for analyzing memory usage. They include PROCMON, TotalView, and Valgrind.

PROCMON is a locally developed application that is ported to all LANL systems and works in parallel. It works on parallel applications. One of its best selling points is that it can monitor other process resource parameters and show graphs correlating memory usage with them, CPU usage, for example. It can be used in two modes; as a passive observer of an unaltered binary or embedded API calls in the applications. In the first mode, it is statistical in nature and can miss local highs or lows, but has very low overhead. In the second mode it will show the usage information at the point the programmer embeds the calls in his application.

TotalView is a parallel debugger that supports all LANL systems. One of its features is that at any breakpoint one can view the process memory usage at that point. This sort of thing could be scripted and the output analyzed later, but it is more convenient feature to use when you already have a good idea about where your problem exists.

Massif is a Valgrind-based tool. It supports LANL's BProc Linux AMD64 and x86 systems, and all relevant compilers. It shows memory usage over time for a process. Valgrind has the ability to work on parallel applications. Massif is a heap profiler. It performs detailed heap profiling by taking regular snapshots of a program's heap. It produces a graph showing heap usage over time, including information about which parts of the program are responsible for the most memory allocations. The graph is supplemented by a text or HTML file that includes more information for determining where the most memory is being allocated. Massif runs programs about 20x slower than normal.

LANL uses Memcheck, Addrcheck, TotalView, Purify, and Insure++ as the primary tools in this area. Memcheck and Addrcheck are Valgrind-based tools which support LANL's BProc Linux AMD64 and x86 systems, and all relevant compilers, as well as the ability to work on parallel applications. Memcheck detects memory-management problems. When a program is run under Memcheck's supervision, all reads and writes of memory are checked, and calls to malloc/new/free/delete are intercepted. As a result, Memcheck can detect most common memory allocation/release and usage errors. Memcheck reports these errors as soon as they occur, giving the source line number at which it occurred, and also a stack trace of the functions called to reach that line. Memcheck tracks addressability at the byte-level, and initialization of values at the bit-level. As a result, it can detect the use of single uninitialized bits, and does not report spurious errors on bitfield operations. Memcheck runs programs about 10 – 30x slower than normal. Addrcheck is a lightweight version of Memcheck which does no uninitialized-value checking. So it detects fewer errors than Memcheck, but programs run about twice as fast as they do on memcheck (about 5 – 20x slower than normal), and a lot less memory is used. Entire KDE sessions have been run under Addrcheck.

Purify is a binary instrumentation product for finding memory corruption and leak detection errors in C/C++ applications. It supports Linux x86, but only for the GCC compiler. Insure++ is an automated run-time C/C++ application testing product that detects elusive errors such as memory corruption, memory leaks, memory allocation errors, variable initialization errors, variable definition conflicts, pointer errors, library errors, logic errors, and algorithmic errors. The source code is instrumented. It supports Linux AMD64 and x86, but only for the GCC compiler.

LANL provided the Understand C/C++ and Understand Fortran products supplied by STI on the production systems to perform static code analysis.

Other products that can be made available for individual use include forcheck and plusFORT.

There are four types of coverage reporting that a tool may provide. As with all tools, if the application is not supported on all ASC platforms it should be architected in such a way that it can be readily ported to a platform of interest.

Line Coverage tells the user which lines have never been executed in the course of the test suite being executed. One may conceptualize this by thinking about source-level debuggers' use of breakpoints. A line coverage tool essentially puts a breakpoint on every

line of code. When the breakpoint is hit, it records that fact.

Branch Coverage tells you whether if has been exercised in ***true*** and ***false***, while or do/while have gone through at least 2 full iterations and exited because the condition became ***false***, for has iterated to completion, and switch/case has reached each case label. *Condition Coverage* tells you whether you have exercised all combinations of conditional statements.

Routine Coverage tells you which subroutines have not been entered. This is macro-level coverage and allows you to design tests that will at least touch all the subroutines if not every line, branch, and branch condition.

LANL has a few tools available for measuring test code coverage. They include Javelina, Intel codecov, and PureCoverage. All of these tools do *Line Coverage* and *Routine Coverage* only.

Javelina allows the user to instrument the binary application and identify the lines of code that are reached during an execution. Javelina depends on Dyninst, a dynamic instrumentation tool. Multiple runs can be combined to show the total coverage for a test suite. It has a graphical capability to show the differences of coverage between two data sets. So, for example, one can easily show what is not tested by a test suite, but is used by a production run of the application or identify code that is not used by any application run (potentially dead code). Work is needed to move to a different instrumentation technique that is available on all the Tri-Lab systems.

Intel's codecov tool has all the same capabilities as Javelina, except that it only works with the Intel compiler on the systems that compiler supports.

PureCoverage is a binary instrumentation product identifies untested code in C/C++ applications. It supports Linux x86, but only for the GCC compiler.

2.3.3.8 Performance Measurement, Analysis, and Tuning

These are tools to give insight into how the applications are functioning and allow the analyst to find prospects for improving the performance while maintaining the correct functionality.

LANL is in the process of installing Open SpeedShop, which will have an MPI message tracing module. Statistical MPI profiling will be enabled by mpiP. Cachegrind is a Valgrind-based tool that shows cache memory usage over time for a process. and has the ability to work on parallel applications. Cachegrind is a cache profiler. It performs detailed simulation of the I1, D1 and L2 caches in the CPU and so can accurately pinpoint the sources of cache misses in your code. It identifies the number of cache misses, memory references and instructions executed for each line of source code, with per-function, per-module and whole-program summaries. It is useful with programs written in any language. Cachegrind runs programs about 20 – 100x slower than normal. The hardware counter profiling tools can give you statistics on cache performance, but not trends of cache usage over time that are available from a cache emulator such as Cachegrind.

LANL has installed PAPI on all LANL systems that support it. Under Linux, the perfctr kernel patch is required (perfctr is not yet a standard part of the Linux kernel).

LANL uses HPCToolkit to profile at higher level constructs such as loops, organize, and make sense of HPM (hardware performance measures) information for an application. PAPI is used as the underlying mechanism for acquiring this data for HPCToolkit on Linux platforms.

Open SpeedShop has a module to profile at higher level constructs such as loops, organize, and make sense of HPC information for an application. It uses PAPI to acquire its data. Open SpeedShop will provide performance profiling for threads and MPI tracing capability.

2.3.4 SNL Linux Software (Chimera)

Sandia Computing produces and supports a Linux software stack (Chimera) for cluster operating environment for HPC Linux clusters. At the core of Chimera is the Red Hat Enterprise Linux 4 (RHEL4) distribution. Some components of the RHEL distribution are modified to meet the demands of high-performance computing installations, operations and support. SNL has added a number of additional cluster-aware components to RHEL.

Chimera contains a set of RPM (Red Hat Package Manager) files, RPM lists for each type of node (compute, management, Gateway, and login), and a methodology for installing and administering clusters. It is produced internally and therefore supports a short list of hardware and software. This approach permits SNL to support a large number of clusters with a single Linux release supported with a small staff and to be agile in planning its content and direction.

In addition to the products of Open Source development described below and the base Red Hat Linux distribution, Chimera includes the following software components:
Kernel — The Linux kernel is based on a Red Hat kernel with additions in the areas of device support for IBA, VFS modifications for Lustre, ECC and FLASH memory device support for Intel motherboard chipsets, miscellaneous bug fixes, and optimized configurations for SNL's hardware.
OpenFabrics — The Open Source IBA software environment used to run parallel programs over the InfiniBand fabric.
Generic hardware monitoring (e.g. lm_sensors, freeipmi) — Hardware monitoring, monitors motherboard chipset sensors such as temperature, fan speed, and power supply voltages.
fping — Fping is a rudimentary node status tool that can ping nodes in parallel. In combination with genders tools, fping can quickly find nodes in the cluster that are turned off or otherwise unreachable on management network.
OpenSSH — OpenSSH provides encrypted remote login/shell service that is PAM-aware.
Fsecure – Fsecure service modified to interact with Sandia Kerberos authentication scheme.
PAM Tools — Pluggable Authentication Module (PAM) for One Time Passwords (OTP) and Kerberos V are used to leverage 's SNL's infrastructure for PAM-aware applications.

Firmware — Firmware images for motherboards, NICs, controllers, and switches, including FLASH/CMOS support software, is included in Chimera. Firmware and support software for power control/serial console hardware is also included.
MPI Test Suite — The Pallas MPI Benchmark (PMB), Effective Bandwidth test (BEFF), and MPI ping-pong test (mping) are packaged with a script to maintain a continuous MPI workload for testing purposes.
NTTCP—The NTTCP TCP bandwidth test is packaged along with genders-aware scripts that can simulate load on the management Ethernet for testing purposes.
Intel Compilers — The Intel EM64T FORTRAN, C, and C++ compilers.
PGI Compilers — The Portland Group Fortran, C, and C++ compilers.
TotalView — The TotalView parallel Debugger from Etnus.
Other Libraries/Tools — Other libraries and tools such as Atlas, Intel MKL (and other Intel libs), Intel Trace, FFTW, ScaLAPAC, OpenGL, VTK, and the Modules package are used to maintain user Linux environment on SNL platforms.

2.3.4.1 SNL Cluster Tools

The following Open Source cluster tools are under active development and various combinations of these tools have been deployed on all SNL Linux clusters:
pdsh — The Parallel Distributed SHell utility executes processes across groups of nodes in parallel. It is also capable of running small MPI jobs on the Elan interconnect.
CAP — The **C**luster **A**dministration **P**ackage (CAP) is an integration, configuration, and systems management toolkit for linux clusters. CAP delivers the right tool sets and methods for managing and maintaining a diverse set of High Performance Computing (HPC) Linux Cluster configurations and architectures. CAP captures these tool sets and methods into component classes.
oneSIS — oneSIS (http://onesis.org) is an open-source software package aimed at simplifying diskless/diskfull cluster management. It is a highly flexible method for deploying and managing a system image for diskless/diskfull systems that can turn any Linux distribution into a master image capable of either being mounted as a read-only NFSroot filesystem, or being deployed onto a node's local disk. One image is sufficient for serving thousands of nodes. Functional groups of nodes are easy to define, and any single node or group of nodes can easily be configured to behave independently.
CHITS module – CIT's cluster hardware issue tracking system (CHITS) used on Cplant, ICC, and NWCC.
Cbench – (http://cbench.sourceforge.net/) Cluster benchmarking, profiling, integration testing, SNL applications, and node diagnostics.
Bohnsack's web interface – Web view of cluster status
modules – Software environment selection/abstraction
Host Monitoring System — TLCC will monitored in-band (while Linux is running) by polling via Ganglia/Supermon (http://ganglia.sourceforge.net and http://supermon.sourceforge.net) . Among the information polled is motherboard sensor information and information about failing hardware devices such as memory and disks. Memory monitoring can be done with the Bluesmoke Linux kernel module (http://bluesmoke.sourceforge.net) reports core CPU or system-level errors. It is mainly concerned with reporting ECC, PCI, machine check, cache, hypertransport, thermal throttling and related events. In addition, IPMI can extract out-of-band monitoring information such as case temperature from some remote power control devices that have this capability. Network statistics (e.g. IBA and ethernet) are collected by inserting

metrics into the Ganglia/Supermon daemons.

2.3.4.2 TORQUE Resource Manager and Job Scheduling

TORQUE (Tera-scale Open-source Resource and QUEue manager) (http://www.supercluster.org/productis/torque/) is a resource manager providing control over batch jobs and distributed compute nodes. It is a community effort and has incorporated significant advances in the areas of scalability, fault tolerance, and feature extensions contributed by NCSA, OSC, USC, the U.S. Dept of Energy, Sandia National Laboratories, PNNL, U of Buffalo, TeraGrid, and many other leading edge HPC organizations.

The primary functions of TORQUE are:
Monitoring the state of nodes in the cluster.
Logically organizes the nodes into partitions with flexible parameters.
Accepting job requests. While TORQUE can support a simple scheduling algorithms, MAUI will manage the more complex scheduling requirements.
MAUI allocates both node and interconnect resources to jobs.
Monitoring the state of running jobs, including resource utilization rates.

TORQUE's APIs support all functions for ease of integration with the MAUI scheduler (http://www.supercluster.org/products/maui/) as well as the parallel job launcher MPIexec (http://www.osc.edu/~pw/mpiexec/).

AIRS is an allocation management system for Linux clusters. Developed by the High Performance Computing Education and Research Center (HPCERC) at the University of New Mexico (UNM) it is an integrated job accounting and user management tool, designed for managing computers that run batch schedulers.

2.3.4.3 Lustre Cluster Wide File System

SNL plans to utilize the Lustre Cluster Wide File System on clusters built up from the TLCC SU. To that end, SNL and Cluster File Systems, Inc., have been actively engaged in developing a production version of Lustre to run over InfiniBand. See http://www.lustre.org/ for more information on Lustre. Currently, Lustre is being deployed on ICC, NWCC, RoSE, and Thunderbird.

2.3.4.4 The Sandia Computing Linux Cluster Support Model

SNL system administrators and software developers work closely with users to resolve problems on production systems. Automated tools such as Nagios (http://www.nagios.org) can be used to generate trouble tickets within the Request Tracker/Asset Tracker (http://atwiki.chaka.net) system in order to capture a complete history of hardware errors on all nodes.

SNL system administrators work closely with the SNL users and Open Source developers and maintainers (Cluster Tools, TORQUE, Maui, Lustre, OpenFabrics) to resolve problems on production systems. Depending on the nature of the package, owners may be the primary developer and fix bugs themselves, or they may be the liaison to an external support resource. External support relationships are primarily developer-to-developer.

2.3.4.5 Integration Testing

Each Chimera release is subject to integration testing that includes regression tests for past problems, basic functionality tests, and real users' applications. Each of the software components is developed asynchronously, but come together in system releases and separate package (TORQUE, Maui, Lustre, Cluster Tools) releases. Due to this separation, system and package testing and installation on production clusters can be scheduled and executed independently. The SNL Linux testbed has a number of small (average 16 nodes) clusters that are available for unit testing of individual software components, integration testing of a complete Chimera release, and debugging of defects that arise in production. These development clusters can accommodate external collaborators working with SNL on the Open Source projects described above.

End of Section 2

3 TLCC Technical Requirements

The end product of the TLCC procurement is a set of highly integrated, well-balanced capacity compute Scalable Units (SU) with at least 10 TF/s, but not more than 25 TF/s as depicted in the SU example in Figure 2-1. Each SU will have compute, IBA networking, gateway, remote partition, and login/service/master resources. These SUs must be combinable in aggregations of at least 1, 2, 4 or 8 SU to form fully functional "capacity" clusters. The Offeror will be responsible for building, passing pre-ship testing with Tri-Laboratory software, delivering, and installing and passing post-ship testing of individual SU. The Offeror, with the receiving Laboratory, will integrate SU into integrated, fully functional clusters and pass cluster acceptance testing. The Offeror will work with the Tri-Laboratory Linux cluster community to integrate necessary device drivers and IBA software into our Linux distributions (see sections 2.3.2, 2.3.3 and 2.3.4). As directed by the University, the Offeror will provide aggregations beyond 2 SU with sufficient additional IBA switches and cables to allow the Tri-Laboratory and the Offeror to construct clusters with full bandwidth, non-blocking IBA interconnects. These SUs shall be capable of supporting a complex workload consisting of small (4-256) medium (257–2,048) and large (2,049-16,384) MPI task count parallel jobs for Tri-Laboratory classified Advanced Simulation and Computing (ASC) and Stockpile Stewardship Program (SSP) simulations. TLCC SUs will reliably run production scientific simulations of a wide number of physical phenomena of importance to all SSP Campaigns and Directed Stockpile Work (DSW). The fully functional SUs and clusters comprised of aggregations of multiple SU must be useful in the sense of being able to deliver a large fraction of peak performance to a diverse scientific and engineering workload. In particular, the SUs and clusters comprised of aggregations of multiple SU must be capable of running a single user application with one MPI task per core over all compute nodes in the cluster. The SUs and clusters comprised of aggregations of multiple SU must also be useful in the sense that the code development and production environments are robust and facilitate the dynamic workload requirements. They must also be easy to install, manage and operate in order to lower the Tri-Laboratory Total Cost of Ownership (TCO).

To satisfy these demanding requirements, we anticipate needing a large set of tightly coupled Linux SU that integrate with Lustre or PanFS global file systems through high-speed external 10 Gb/s Ethernet networking. Our requirement is to have these SU built from commodity AMD x86-64 or Intel EM64T (or binary equivalent) nodes containing at least four (4) microprocessor sockets. These SUs shall have IBA 4x DDR (or faster) compatible interconnect consisting of IBA switches, cables, and adapters. In addition, these SU shall have 1 and 10 Gb/s Ethernet connections for external networking.

This procurement is structured to span two US Government Fiscal Years (GFY) from third Quarter Calendar Year 2007 (3QCY07) through 3QCY08. During this period of time, there will be several advancements in COTS technology utilized in any proposed SU configuration. As such, the Offeror shall provide these technology enhancements to the Tri-Laboratory community in future deliveries of SU, but need not upgrade past delivered SU.

Mandatory Requirements (designated MR) in this Statement of Work are performance features that are essential to Tri-Laboratory requirements, and an Offeror must satisfactorily propose all Mandatory Requirements in order to have its proposal considered responsive.

Target Requirements (designated TR-1, TR-2, or TR-3) in this Statement of Work, are features, components, performance characteristics, or other properties that are important to the Tri-Laboratory, but will not result in a nonresponsive determination if omitted from a proposal. Target Requirements add value to a proposal. Target Requirements are prioritized by dash number. TR-1 is most desirable to the University, while TR-2 is more desirable than TR-3. Target Requirement responses will be considered as part of the proposal evaluation process.

In addition to MRs and TRs identified in this Statement of Work, the Offeror may choose to propose any additional features (i.e., Offeror proposed features) consistent with the objectives of the TLCC procurement and the Offeror's project plan, which the Offeror believes will be mutual benefit. MRs, TRs, and additional features proposed by the successful Offeror, and of value to the Tri-Laboratory, will be stated as firm requirements in a final negotiated Statement of Work and incorporated in the resulting TLCC Subcontract.

3.1 High-Level Hardware Summary (TR-1)

Offeror will provide a high-level overview of the proposed SU design (section 3.1.1) and its evolution (section 3.1.3) over the 3QCY07 through 3QCY08 timeframe. The intent of this section is to have in one place a technical summary of the Offeror's proposed SU deliveries. It is vital that the Offeror make absolutely clear in the response to these subsections, what will be delivered and when.

3.1.1 SU High-Level Architecture

Offeror's response to this section will contain a detailed description of the proposed TLCC SU and the proposed evolution of this SU technology over time. The features and functionality of all major components of the SU shall be discussed in detail. The Offeror will provide an architectural diagram of the TLCC SU, similar to Figure 2-1, labeling all component elements and providing bandwidth and latency characteristics (speeds and feeds) of and between elements. The Offeror will provide an architectural block diagram for each TLCC node type bid, labeling all component elements and providing bandwidth and latency characteristics (speeds and feeds) of and between elements. The node architectural diagrams will specifically show and label the chipset used and denote independent PCIe buses and slots and label these with bus widths and speeds. The Offeror will provide an architectural block diagram of the proposed IBA interconnect for the SU and for combining SU in at least 1, 2, 4 and 8 multiples similar to Figure 2-2. Offeror will provide a rack layout diagram for the proposed SU similar to Figure 2-4 and floor layout for at least four clusters consisting of aggregations of four SU each, similar to Figure 2-5. If Offeror proposes to deliver different SU packaging configurations with differing rack layouts in order to meet site specific power, cooling requirements (see section 3.2.12 and subsections), then a rack layout diagram for each proposed SU packaging configuration will be provided. Any alternative cooling strategies with non-trivial facilities impacts should be described.

3.1.2 SU Requirements Summary Matrix

The following matrix identifies the highest priority technical requirements (TR-1) and will be completed in its entirety. Entries shall be labeled N/A if the requirement is not offered. In addition, the system requirements summary matrix will be completed for any alternate proposed systems submitted.

Index	Requirement Description	Qty	Offeror Response
	Compute node product designation		
2.	Compute node form factor	NA	
3.	Compute node processor type, speed, and cache sizes		
4.	Compute node memory bus type and speed		
5.	Compute node chip set designation		
6.	Compute node number of expansion busses and types		
7.	Compute node number and type of expansion bus slots for each bus		
8.	Type and size of compute node memory		
9.	Compute node blade-chassis type and configuration, if applicable		
10.	LSM node product designation		
11.	LSM node processor type, speed and cache sizes		
12.	LSM node memory bus type and speed		
13.	LSM node chip set designation		
14.	LSM node number of expansion busses and types		
15.	LSM node number and type of expansion bus slots for each bus		
16.	Type and size of LSM node memory		
17.	Type and size of LSM node local SATA disk		
18.	Gateway node product designation		
19.	Gateway node processor type, speed and cache sizes		
20.	Gateway node memory bus type and speed		
21.	Gateway node chip set designation		
22.	Gateway node number of expansion busses and types		

Index	Requirement Description	Qty	Offeror Response
23.	Gateway node number and type of expansion bus slots for each bus		
24.	Number and type of each PCIe expansion card(s) installed in each Gateway		
25.	Type and size of gateway node memory		
26.	RPS node product designation		
27.	RPS node processor type, speed and cache sizes		
28.	RPS node memory bus type and speed		
29.	RPS node chip set designation		
30.	RPS node number of expansion busses and types		
31.	RPS node number and type of expansion bus slots for each bus		
32.	Type and size of RPS node memory		
33.	Type and size of RPS SATA disks and RAID5 config. Indicate SATA RAID5 packaging solution (e.g., internal to node, external expansion chassis)		
34.	Number and type of each PCIe expansion card(s) installed in each Gateway		
35.	RAID5 controller designation and interface types and numbers		

3.1.3 SU Evolution Overview (TR-1)

The Tri-Laboratory requires that the SU that are aggregated into a specific cluster at any site be as identical as possible. However, we also require that when processor, interconnects (networking), memory and disk technology elements advance during the lifetime of the contract resulting from this procurement, that these enhancements will be integrated into future SU deliveries without perturbing SU cost or reliability significantly. Offeror need not propose to upgrade SU hardware after delivery. Offeror will offer at least the following technology enhancements:

1. Processor frequency improvements within the same cost and power envelops
2. New processor socket and/or chipset improvements
3. New processor cores

4. SATA disks with higher capacity
5. Higher speed and capacity memory improvements
6. Interconnect bandwidth and latency improvements

Offeror provide at least the following information for these technology improvements. Overall SU Impact should be rated as one of (low, medium or high) and then major components that are impacted will be listed. Offeror will use "low" impact designation to indicate that other no other major components are impacted by the change. Offeror will use "medium" impact designation to indicate that other major components of the SU require update, but not a new design and the SU architecture does not change substantially. Offeror will use "high" impact designation to indicate that other major components of the SU require redesign and/or the SU architecture does changes substantially.

Item	Item Upgrade	Delivery Qtr	Attribute	Overall SU Impact
Processor	Speed Bump	3QCY07	X.X GHz clock	Low
		4QCY07	X.X GHz clock	Low
		1QCY08	X.X GHz clock	Low
		2QCY08	X.X GHz clock	Low
		3QCY08	X.X GHz clock	Low
Processor	New socket		Socket Type, clock	Medium, new motherboard, new memory type/speed
Processor	Next Generation Processor	3QCY07	Processor name, socket, GHz clock, power	High, new motherboard, node design, new memory type/speed, new node design
Processor	OTHER			
Memory	DDR2 or FBD	3QCY07	More Bandwidth	Medium, new motherboard
Memory	OTHER			
IBA	4x DDR → 4x QDR	3QCY08	4GB/s bandwidth	Medium, new switches
IBA	Other			
Local Disk	HDD capacity bump	3QCY07	800 GB	Low
		4QCY07	1 TB	Low
Local Disk	OTHER			

For proposed technology improvements that have medium or high impact to SU architecture design, Offeror will provide high-level SU architectural diagrams defined in section 3.1.1 for each.

3.2 SU Hardware Requirements

For each of the following SU hardware requirements, provide information for the first SU insanitation only. Changes to the proposed SU hardware to meet these requirements over the contract timeframe are covered in section 3.1.3.

3.2.1 TLCC SMP Scalable Unit (MR)

Each SU the Offeror provides shall be based on at least two socket AMD64 or Intel EM64T (or binary compatible) nodes. All nodes in SUs to be aggregated into a specific cluster at any site shall be of the same processor and chipset revision. There will be four node types on the IBA switch: compute nodes; gateway nodes; login/service/master node, remote partition server. All nodes shall be attached to the management Ethernet network. The login/service/master node shall attach to a 10 Gb/s Fiber Ethernet infrastructure. The gateway nodes shall attach to the 10 Gb/s Ethernet infrastructure through two 10 Gb/s Ethernet Fiber connections. The login/service/master nodes shall also attach to a 1 Gb/s Fiber Ethernet infrastructure. The remote partition server shall be attached to the management Ethernet with 1 Gb/s copper Ethernet and also have 1 Gb/s fiber Ethernet infrastructure..

3.2.2 TLCC SMP Four Socket Configuration (TR-1)

The SU compute and login/service/master nodes the Offeror provides will be based on at least four socket AMD64 or Intel EM64T (or binary compatible) nodes. If required to meet the node performance requirements below, the gateway and remote partition server nodes may also be four socket.

3.2.3 SU Peak (TR-1)

Each SU the Offeror provides will have a peak 64b floating point processing power of at least 10.0 teraFLOP/s ($10x10^{12}$) and not more than 25.0 teraFLOP/s ($2.5x10^{13}$). In the following sections "F" in the B:F ratios denotes the Offeror's proposed node or SU 64b floating point peak rate.

3.2.4 Number of SU (MR)

N, the number of SU required over the lifetime of the contract is 365 teraFLOP/s divided by the SU delivered in 3QCY2007 peak rounded to the nearest integer.
Example: a 20.3 teraFLOP/s peak SU yields N = round(365/20.3) = 18. An 12.3 teraFLOP/s peak SU yields N = round(365/12.3) = 30.

3.2.5 TLCC Node Requirements

The following requirements apply to all node types except where superseded in subsequent sections.

3.2.5.1 Processor and Cache (TR-1)

The SU nodes will be configured with at least four AMD64 or Intel EM64T (or binary compatible) quad core processor dies with at least 1 MB of L2 or L3 on-chip cache. Each socket will utilize less than 100 Watts when running a copy of Linpack on every core and one speed grade slower than the fastest available in that wattage category available when the SU is built.

3.2.5.2 Chipset and Memory Interface (TR-2)

The SU nodes will be configured with a chipset compatible with DDR2-5400(667 MHz) or faster SDRAM delivering a peak of 10.7 GB/s memory bandwidth per processor socket from a total of 4 DIMM memory slots per processor socket. If B represents the delivered aggregate peak memory bandwidth, then this represents a memory bandwidth B:F of at least 0.25. The SU memory configuration will be architected and configured to deliver low memory latency and high bandwidth to ASC applications. The node will be configured with at least two PCIe 8x (or faster) busses with one slot per bus. It is highly preferred that the chipset for all nodes in the cluster be the same.

3.2.5.3 Node Delivered Memory Performance (TR-1)

The SU nodes will be configured to deliver at least 5.4 GB/s per processor socket of memory bandwidth when running one copy of the streams benchmark per core. Offeror will report with proposal the delivered streams performance running one copy of the streams benchmark per core, for each bid TLCC node type. See http://www.llnl.gov/asci/purple/benchmarks/ or http://www.cs.virginia.edu/stream/ for the streams benchmark.

3.2.5.4 Node I/O Configuration (TR-2)

The SU nodes chip set will be configured with two PCIe 8x (or faster) buses and two PCIe 8x slots, for 10 Gb/s Ethernet, and at least one PCIe 8x (or faster) bus and PCIe 8x slot or one HyperTransport™ bus and HTX slot for the IBA HCA. The SU node PCIe and HyperTransport™ infrastructure will be fully compatible with the proposed IBA HCA.

3.2.5.5 Node Memory (TR-1)

The SU nodes will be configured with at least 2.0 GB memory per processor core. The memory should be DDR2-5400 (667 MHz) or faster SDRAM registered DIMMs with ECC and chipkill. This would result in 32.0 GB (16x2 GB DIMMs) for 4-socket 4-core processor solutions and 16.0 GB for 2-socket 4-core processor solutions.

3.2.5.6 Node Power (TR-2)

The SU nodes components will utilize less than 200 Watts per socket plus 150 Watts when running a copy of Linpack on every core and one speed grade slower than the fastest available in that wattage category available when the SU is built. Power supplies will have a minimum efficiency of 80% when tested at each of the following three load conditions: 20%, 50% and 100% of rated power supply output. The test procedure for such measurements may be found at http://www.efficientpowersupplies.org . Power supplies will maintain a true power factor of 0.9 or greater at 100% of rated power supply output. Unrated power supplies (of any output) can be sent to http://80plus.org and rate for a fee (~ $400).

3.2.5.7 Linux Access to Memory Error Detection (TR-1)

The SU nodes will include a hardware mechanism to detect correctable and uncorrectable memory errors from Linux. This hardware mechanism will be capable of sending a non-maskable interrupt (NMI) or machine check exception when an uncorrectable error occurs, so that software may take immediate action. When a correctable or uncorrectable memory error occurs, this hardware mechanism will provide sufficient information so

that software may identify the affected failed or failing DIMM FRU (i.e. the exact DIMM location on the motherboard). Node memory controller will expose defective memory modules that the chipkill or other functionality compensates for to the operating system down to the DIMM FRU level. Documentation that is sufficiently well written and complete so that Tri-Laboratory personnel can actually program this hardware facility will be delivered with the SU.

3.2.5.8 No Local Hard Disk (TR-1)

The SU nodes will be configured **without** a hard disk drive.

3.2.5.9 Node Form Factor (TR-2)

The Offeror will provide SU nodes with equivalent form factor of not more than 2U. Denser (including blades) solutions, meeting the power and cooling requirements in section 3.2.12 and facilities requirements in section 5 are desired.

3.2.5.10 Integrated Management Ethernet (TR-1)

The TLCC SU nodes will have at least one integrated copper 100 BaseT or faster (e.g., 1000 Base-TX (Copper)) management Ethernet.

3.2.5.11 Node BIOS (MR)

Offeror shall provide all nodes of an SU shall with a fully functional BIOS that shall take a node from power on or reset state to the start of the Linux kernel as loaded from a network connection or local disk (if a local hard disk is installed). The requirements below apply to all nodes types except where called out as specific to a particular node type.

3.2.5.11.1 Node BIOS Type Options (MR)

One of two options shall be employed to meet the TLCC node BIOS requirements. All other requirements of this BIOS section, as well as other requirements throughout this document, shall apply to either BIOS version.

3.2.5.11.1.1 LinuxBIOS (TR-1)

Offeror may install LinuxBIOS on all nodes, in their standard configurations along with a complete set of tools to manage LinuxBIOS (see http://www.linuxbios.org). Offeror may provide complete LinuxBIOS documentation for the delivered SU nodes. This documentation may include description of all parameters, default factory settings and how to modify them, use and abuse of the tools to manage LinuxBIOS. Offeror may contribute the LinuxBIOS modifications for the compute and gateway nodes back to the LinuxBIOS Open Source Community.

3.2.5.11.1.2 Standard BIOS (TR-2)

Offeror may install a proprietary, Offeror supported BIOS on all nodes, in their standard configurations along with a complete set of tools to manage BIOS. Offeror may provide complete BIOS documentation for the delivered SU nodes. This documentation may include description of all parameters, default factory settings and how to modify them, use and abuse of the tools to manage BIOS.

3.2.5.11.2 Remote Network Boot (MR)

All provided node BIOS shall be capable of booting the node from a remote OS image across the Ethernet Management network. Booting over Ethernet shall utilize PXE, BOOTP, DHCP, or other public protocols.

3.2.5.11.3 Remote IBA Network Boot (TR-1)

All provided node BIOS may be capable of booting the node from a remote OS image across the IBA interconnect network Booting over IBA may utilize an IBA standard protocol like PXE, SRP, iSER, IPoIB, or similar.

3.2.5.11.4 Node Initialization (TR-1)

The node BIOS initialization process may complete without human intervention (e.g., no F1 keystroke to continue) or fail with an error message written to the console. The time required for a node's BIOS to take the node from a power on state (or reset state) to the start of the loading of the Linux kernel boot image may be less than thirty (30) seconds. Note that this includes the POST phase configured with minimal POST hardware checks. Shorter times are clearly desirable and achievable.

3.2.5.11.5 Error/Interrupt Handling (MR)

The delivered BIOS shall not block reporting, interrupts or traps from chipset errors, memory errors, sensor conditions, etc. up to the Linux Kernel level. All of these conditions shall be passed to the Linux Kernel level for handling. The delivered BIOS shall not attempt to respond these conditions directly other than a request to reboot or reset the node.

3.2.5.11.6 BIOS Security (MR)

The principal function of the provided BIOS shall be to perform node hardware initialization, power-on self testing and turn over operation to the OS boot loader. The BIOS shall not perform any kind of extraneous automated or uncontrolled I/O to disks, networks, or other devices beyond that required to read or write BIOS images, CMOS parameters, and device or error registers as required for booting, BIOS configuration, power-on testing, and hardware diagnostics and configuration. Under no circumstances shall the BIOS itself be allowed to directly capture or write user data to any location. If the BIOS has such capabilities but is configured to disable them, then a formal testing process and the results of such tests shall be provided which conclusively demonstrate such features are in fact disabled.

3.2.5.11.7 Plans and Process for Needed BIOS Updates (MR)

A written plan shall be submitted with bid proposal(s) outlining Offeror's plans and process to provide BIOS updates to address problems or deficiencies in the areas of functionality, performance, and security. The plan shall outline a process which Offeror shall follow to identify, prioritize and implement BIOS updates in general and for addressing any specific Tri-Lab and/or DOE security related issues and concerns that may be raised. The Offeror's plan shall be finalized after contract award and included as an early deliverable the Statement of Work for the contract.

3.2.5.11.8 Failsafe/Fallback Boot Image (TR-1)

The BIOS may employ a failsafe/fallback booting capability in case of errors in the default BIOS image, during the default BIOS boot process, or during a BIOS flash operation. The node may reach a minimum state allowing for the diagnosis of errors. This capability may not require the use of any external writeable media (e.g. floppy, flash disk, etc.) and preferably no external media at all.

3.2.5.11.9 Failsafe CMOS Parameters (TR-1)

The BIOS may have failsafe default CMOS parameter settings so that the serial console interface will function if the CMOS values are unintentionally reset to the default values (e.g., the CMOS battery fails). The failsafe parameters may be modifiable in accordance with 3.2.5.11.13 below, if provided.

3.2.5.11.10 Serial Console after Failsafe/Fallback Booting (TR-1)

After a failsafe/fallback BIOS condition (section 3.2.5.11.8) the BIOS would enable at a minimum a serial console which can operate with the remote management capabilities as specified in section 3.2.7 to allow remote access and management of the node. This capability may not require the use of any external media of any kind (CD-ROM, DVD, floppy, flash device, etc.).

3.2.5.11.11 BIOS Upgrade and Restore (TR-1)

Offeror may provide a hardware and software solution that allows a Linux command line utility or utilities to update (flash) or restore the BIOS image(s) in the flash BIOS chip and to verify the flash image(s). This mechanism may not require booting an alternative operating system (e.g. DOS).

3.2.5.11.12 CMOS Parameter Manipulation (TR-1)

The Offeror may provide a mechanism to read (get) and write (set) individual CMOS parameters from the Linux command line. In particular, CMOS parameters such as boot order may be modifiable from Linux command line. In addition, there may be accessible CMOS parameters that can disable any power management or processor throttling features of the node. The Linux command line utility may also allow reading all CMOS parameters (backup) and writing all CMOS parameters (restore). Reading or writing CMOS parameters may not require booting an alternative operating system or interacting with BIOS menus or a BIOS command line interface. Such a capability may function from within Linux command scripts and also allow the use with the Linux Expect utility and not require navigation of BIOS menus.

3.2.5.11.13 BIOS Command Line Interface (TR-2)

A Linux shell command-line interface may be provided to interact with the delivered BIOS. It provides access to any other BIOS functionality above and beyond that described in section 3.2.5.11.12 and may be integrated with the tools associated with CMOS manipulation if both are provided. This capability may function from within Linux shell scripts and also may allow the use with the Linux Expect utility and may not require navigation of BIOS menus.

3.2.5.11.14 Serial Console over Ethernet Support (TR-2)

The BIOS may directly provide a remote serial console capability over an Ethernet management network or over the IBA interconnect instead of over an actual serial port. This should support the remote interactive management features of section 3.2.7.

3.2.5.11.15 Power-On Self Test (TR-2)

The POST may be comprehensive, detect hardware failures and identify the failing FRU during the power up boot sequence. All POST failures should be reported to the serial console and through the remote management solution in section 3.3.7.

3.2.5.11.16 BIOS Security Verification (MR)

Offeror shall propose either or both of the two following requirements in order to verify that the provided BIOS meets the security related requirements (section 3.2.5.11.6).

3.2.5.11.16.1 BIOS Buildable Source Code (TR-1)

Offeror may provide licenses and actual buildable source code for delivered node BIOS and all future updates. The buildable BIOS source may be provided which allows the Tri-Lab community to build and flash a usable BIOS and failsafe BIOS. Buildable source code may also be provided for all necessary Linux kernel modules, BIOS and CMOS utilities. A source license may be Tri-Lab-wide and may allow this community to work together by exchanging information, source code patches and source code.

3.2.5.11.16.2 BIOS Security Certification (TR-1)

The BIOS may be provided with a vendor certification document indicating that the BIOS strictly meets the conditions specified in section 3.2.5.11.6. This certification document may become part of a DOE Security Plan for classified operations of these systems at the Tri-Lab sites. Future updates of the BIOS must be accompanied by a new certification appropriate for that specific BIOS version.

3.2.5.12 Programmable LED(s) (TR-3)

Offeror will provide nodes with either programmable LED(s) or a programmable front Alpha Numeric panel for run-time diagnostics. Access to these may be made available to the BIOS as well as Linux.

3.2.6 IBA Interconnect (MR)

The SU shall be built with an InfiniBand Architecture (IBA) compatible interconnect (www.infinibandta.org/specs version 1.1). A network is IBA compatible if it is capable of running the OpenFabrics software stack, section 3.3.1, 3.3.2) with a port of the HCA drivers.

3.2.6.1 Node HCA Functionality (TR-1)

The SU IBA network will have at least one IBA 4x DDR (or faster) with PCIe 8x or HyperTransport™ (or faster) HCA in every node. The HCA will not have local memory on the card (mem-free).

3.2.6.2 Node Bandwidth, Latency and Throughput (TR-1)

The SU network will have at least one IBA 4x DDR (or faster) with PCIe 8x or HyperTransport™ (or faster) HCA. The SU IBA network shall deliver node MPI link bi-directional bandwidth between all node pair combinations B:F of at least 0.028 as measured with one MPI task per core on all node pairs. The SU IBA network shall deliver an MPI ping/pong latency (round trip divided by two) of no more than 1.7 µs as measured between any and all two MPI task pairs in the SU, each with one MPI task per core. The SU IBA network shall deliver an aggregate processing rate of at least $2.7x10^7$ messages per second utilizing an application with one MPI task per core on each node, between any two nodes in the SU. Subcontractor will report below the compute node delivered MPI bandwidth, latency and messaging throughput benchmarks over IBA 4x DDR (or faster) HCA attached to PCIe 8x or HyperTransport™ (or faster) bus between two nodes utilizing 1 MPI task/node, 1 MPI task/socket and 1 MPI task/core on each node.

We recommends the "com" benchmark from the ASC Purple Presta MPI Stress Test suite. Presta can be obtained from:
http://www.llnl.gov/asci/purple/benchmarks/limited/presta/

3.2.6.3 Fully Functional IBA Topology (TR-1)

The Offeror will propose a fully functional IBA 4x DDR (or faster) interconnect with all required hardware (e.g., switches, cables, connectors, and HCA) and software stack (section 3.3.1, 3.3.2). The delivered IBA network will be capable of running OFED V1.2 or later. The delivered IBA network hardware and software will be capable of driving at least one type of copper to optical cable transducer through the copper CX4 interface. Offeror will assist the OpenFabrics community on porting the OFED V1.2 stack to the SU interconnects. For SU destined for 2 SU clusters, the SU network will have at least one IBA 4x DDR (or faster) interconnect fabric with fat-tree topology and full, non-blocking bisection bandwidth that allows all nodes in the 2 SU cluster to communicate with every other node in the 2 SU cluster. For SU destined for 4, 6 or 8 SU clusters, the SU network will have at least one IBA 4x DDR (or faster) interconnect fabric with fat-tree topology and full, non-blocking bisection bandwidth that allows all nodes in the SU to communicate with every other node in the 4, 6 or 8 SU, possibly through the spine switch. The SU network will have as many ports available for SU nodes as available for connection to spine switches in a full, non-blocking bandwidth configuration. Larger switches (288-port or larger) are desired as the spine switches, with smaller (24-port or larger single switch ASIC) switches are desired for the leaf switches for the 4, 6 and 8 SU cluster configurations.

3.2.6.4 IBA Interconnect Reliability (TR-1)

In order to have 8 SU cluster aggregations of SUs that can support ASC production workloads, the Offeror will propose a IBA interconnect infrastructure that is reliable in the sense that it meets or exceeds the following:

* HCA failure requiring replacement is less than one per two months per 1,152 nodes. This corresponds to an HCA MTBF of over 1.558 million hours.
* HCA failures due to catastrophic state or reset or becomes non-responsive less than one per two months per 1,152 nodes.
* Switch FRU failure rate less than one per two month per 1,152 node cluster. This corresponds to a switch FRU MTBF of over 473 thousand hours.

Less than one link loss per quarter per 3,456 links. This corresponds to a link loss of over 7.4 million hours.
* Links drops back to 1x SDR or 1x DDR or 4x SDR no more than once per 1,024 restarts of Open SM per 1,152 links.
* Link Transmission Errors (any data corruption corrected or not loss in any packet transmitted or loss of packets that cause packet timeout with default timeout of $4*2^{12}$ µs)

3.2.6.5 Multi-SU Spine Switches (TR-1)

The SU network for SU destined for 2 SU cluters will be capable of directly connecting two SU together without additional switches. The SU network will be capable of expanding to a cluster with at least 4, 6 or 8 SU with the addition of Offeror supplied spine switches and cables.

3.2.7 Remote Manageability (TR-1)

All nodes will be 100% remotely manageable, and all routine administration tasks automatable in a manner that scales up to a cluster aggregation of eight SU. In particular, all service operations under any circumstance on any node must be accomplished without the attachment of keyboard, video, monitor, and mouse (KVM). The Tri-Laboratory community intends to use the open source tools PowerMan and ConMan for remote management, and therefore the Offeror will propose hardware and software that is reasonably compatible with this software environment and provide any software components needed to integrate the proposed hardware with these tools. Areas of particular concern include remote console, remote power control, system monitoring, and system BIOS and LinuxBIOS.

The Offeror will fully describe all remote manageability features, protocols, APIs, utilities and management of all node types bid. Any available manuals (or URLs pointing to those manuals) describing node management procedures for each node type will be provided with the proposal.

All remote management protocols, including power control, console redirection, system monitoring, and management network functions must be simultaneously available. Access to all system functions within the SU, must be made available at the Linux level so as to enable complete system health verification.

Offeror will propose a remote management solution that is based on section 3.2.7.1 or section 3.2.7.2.4, but not both. The Tri-Laboratory community has a strong preference for reliable, complete solutions based on section 3.2.7.1.

3.2.7.1 Traditional Remote Management Solution (TR-1)

A “traditional” remote management solution (TRMS) is preferred. However, if the proposed SU configuration can’t implement TRMS, then a fully functional IPMI 2.0 and BMC (section 3.2.7.2) style remote management solution is an acceptable alternative.

3.2.7.1.1 Remote Console (TR-3)

The TRMS will interface to the console port of on every node in the SU. The TRMS console interfaces will be aggregated in Terminal servers on the Management Ethernet. The TRMS will interface to the power plug on every node of the SU. The

TRMS power interfaces will be aggregated in a power control device on the Management Ethernet.

3.2.7.1.2 Remote Node Power Control over Management LAN (TR-3)

Remote access to power control device over the management LAN will be accomplished through a command line interface that is can easily be scripted with the Linux Expect utility. The power control device will be capable of turning each node's power off and turning each node's power on and querering the power state of the node. The power control infrastructure will be able to reliably power up/down all nodes in the SU simultaneously. Reliable here means that 1,000,000 power state change commands will complete with at most one failure to actually change the power state of the target nodes.

3.2.7.2 IPMI and BMC Remote Management Solution (TR-1)

Node remote management will be accomplished with IPMI 2.0 and a baseboard management controller (BMC). In the event that the BMC is not integral to the base motherboard, the BMC daughter card (or equivalent) will be proposed. The Offeror will provide a fully compliant IPMI 2.0 implementation. Any security relevant features in the IPMI specification must be supported and configurable. All IPMI functions will be utilized from Linux and there should be no requirements for any DOS based utilities. Offeror will provide a Linux command line utility or utilities that allow upgrade and verification of BMC firmware and BMC configuration values. The command line utility will allow reading of necessary BMC configuration parameters and writing of necessary BMC configuration parameters. Linux command line utilities for firmware upgrades must be able to work in-band. An out-of-band only firmware upgrade solution is not acceptable. BMC configuration will be based on FreeIPMI bmc-config utility (http://www.gnu.org/software/freeipmi). Offeror's proposed solution will not require OEM IPMI 2.0 extensions for remote manageability. Offeror will publicly release documentation on any OEM specific system event log or IPMI error codes. All power supply, processor state, and sensors listed in the Sensor Type Codes table (Table 42-3 of the IPMI 2.0 Specification) will supply the sensor values corresponding to those given in that table. The Offeror may not provide their own sensor values and interpretations.

The BMC and motherboard hardware will provide the following sensor data through IPMI 2.0:
Each and every fan within the node
Temperature of every processor die
All motherboard temperature sensors
Voltage supply to each socket
Processor state
Power supply state

Temperature sensors should be designed to be insensitive to manufacturing tolerances, e.g. CPU thermal diode readings should utilize the dual-sourcing current or more accurate methodology. Offeror will publicly release documentation on any OEM specific motherboard sensors so that the sensors can be interpreted correctly. In particular, sensor accuracy and precision should be stated for each sensor. An individual sensor will be

provided for each power supply and processor (or processor core) that exists in the system. A single sensor that represents multiple power supplies or processors/cores is not acceptable.

The IPMI solution will allow the following requirements below to be met concurrently over the SU management LAN.

3.2.7.2.1 ConMan Access to Console via Serial over LAN (TR-1)

ConMan will access all node consoles simultaneously via the IPMI 2.0 Serial Over LAN (SOL) for serial console access. The SOL implementation will meet requirements for serial console listed in Sections 3.2.7.3.1 through 3.2.7.3.3.

3.2.7.2.2 LAN PowerMan Access (TR-1)

PowerMan will access the BMC power management features on every node in the SU simultaneously via the freeIPMI ipmipower tool. The BMC power management features will be capable of turning each node's power off and turning each node's power on. The BMC based power control infrastructure will be able to reliably power up/down all nodes in the SU simultaneously. Reliable here means that 1,000,000 power state change commands will complete with at most one failure to actually change the power state of the target nodes.

3.2.7.2.3 LAN Management Access (TR-1)

All other node management functions will be accomplished via a remote mechanism to every node in the SU simultaneously. The remote node management mechanism implementation will never allow message buffer overflow or data corruption conditions.

3.2.7.2.4 IPMI User Role Division (TR-2)

An IPMI command (or commands) will be provided that will allow the University to configure each BMC user to use only a configured set of IPMI commands. For example, one IPMI user may be configured to only be able to use those IPMI commands to access sensor readings. Another IPMI user may be limited to only perform remote power control commands.

3.2.7.2.5 Traditional Remote Management Backup Plan (TR-2)

Offeror will propose a traditional remote management backup plan meeting the requirements in section 3.2.7.1 as a back up plan should the IPMI 2.0 based solution prove to be unworkable and/or unreliable.

3.2.7.2.6 Additional IPMI Security Requirements (MR or TR-1?)

Due to security policies in place at the University, several additional IPMI features not in the IPMI specification must be provided by the Offeror so that security policies can be met. The additional security features will be provided via IPMI commands and sensor events which are capable of being executed and read via the FreeIPMI. IPMI commands and sensor events will be available to be published in the GPL software released by the University. The Offeror will work with the University to get these or similar changes into the IPMI specification. Binary or web based tools which supply the features are not acceptable.

3.2.7.2.6.1 Disabling a User After Failed Out-of-band Authentication Attempts (TR-1)

Each BMC user will be disabled after a configurable number of out-of-band authentication failures has been reached. An IPMI command will be provided that will allow the University to configure each MBC user's limit. The method by which a BMC user is re-enabled after lockout may be through the IPMI Set User Password command or an additional IPMI command provided by the Offeror. The commands offered will be usable via FreeIPMI's libfreeipmi and ipmi-raw. The commands will be included in FreeIPMI's bmc-config tool at a later time. The failed authentication attempts will cover both IPMI 1.5 and IPMI 2.0 authentication attempts.

3.2.7.2.6.2 Failed IPMI Out-of-band Authentication Attempts (TR-1)

An IPMI sensor event will be provided that indicates if a BMC user had a recent failed out-of-band authentication attempt. The event will be readable via FreeIPMI's ipmi-sel and will be available for reading by platform event filtering tools. The failed authentication attempts will cover both IPMI 1.5 and IPMI 2.0 connection attempts.

Note that the "Platform Security Violation Attempt" sensor in Table 42-3 of the IPMI 2.0 specification is not acceptable. It does not provide a mechanism to inform the University which user the failed password attempt occurred with and semantically only covers failed IPMI 1.5 authentication attempts.

3.2.7.3 Remote Management Solution Requirements (TR-1)

The following requirements apply to both the IPMI 2.0 (section 3.2.7.2) and TRMS (section 3.2.7.1) solutions.

3.2.7.3.1 Serial Console Redirection (TR-1)

All BIOS interactions will be through a serial console. There will be no system management operations on a node that require a graphics subsystem, KVM or DVDROM or floppy to be plugged in. In particular, the serial console will display POST messages including failure codes, operate even upon failure of CMOS battery and provide for a mechanism to remotely access the BIOS.

3.2.7.3.2 Dedicated Serial Console Communications (TR-2)

The serial console communication channel on every node will be available simultaneously for console logging and interactive use at all times. This is to ensure that all console output is logged and Expect scripts that perform console or service processor actions do not interfere with each other or with console logging. All console output must be available for logging at all times with no dropped or corrupted data. The serial console will operate after node crash or hang. The serial console will operate during a network boot.

3.2.7.3.3 Serial Console Efficiency (TR-1)

The serial console communication channel will support a baud rate of 115,200 or greater. If an IPMI 2.0 solution is offered, the BMC will transfer AES-128 encrypted

character data at a rate equivalent to a traditional 115200 baud serial console.

3.2.7.3.4 Flow Control (TR-1)

Flow control will **not** be required for serial console communication.

3.2.7.3.5 Peripheral Device Firmware (TR-2)

Offeror will provide Linux a utility or utilities for saving, restoring, verifying (including printing version number) firmware for any peripheral devices supplied.

3.2.7.3.6 Remote Network Boot Mechanism (TR-1)

The node BIOS will support booting an executable image over the management Ethernet utilizing PXE, BOOTP or DHCP with the vendor BIOS and LinuxBIOS. Console data and power management functionality must be available during network boot process.

3.2.8 Gateway Node Requirements (TR-1)

The following Requirements are specific to the gateway nodes and supersede the general node requirements above.

3.2.8.1 Gateway Node Count (TR-1)

The Offeror will configure the SU with six (6) Gateway nodes for SU peak between 10 teraFLOP/s and 20 teraFLOP/s and twelve (12) Gateway nodes for SU peak above 20 teraFLOP/s.

3.2.8.2 Gateway Node Configuration (TR-1)

The SU gateway nodes will be used for file system and other IP based connectivity between the compute nodes and the global file system and other IP based communications infrastructure. For Lustre, the gateway will run the "LNet Router code," that routes LNet/Verbs/IBA to LNet/IP/10 Gb/s Ethernet. For PanFS, the gateway will route IP packets between IP/IBA to IP/10 Gb/s Ethernet. Offeror will configure the SU with a minimum number of gateway nodes to achieve a delivered gateway bandwidth throughput of 0.00035 B:F, where B is the aggregate number of Bytes/s that the gateways can route IP packets between the IBA and 10 Gb/s Ethernet networks. For a 20 teraFLOP/s this requirement translates into an aggregate gateway delivered IP routing bandwidth of 7.2 GB/s. The Offeror's gateway should carefully balance the IBA HCA and 10 Gb/s Ethernet network *delivered* IP bandwidths. Gateway node may also include at least one 1 Gb/s 1000Base-TX Ethernet interface in addition to any management Ethernet. This interface will be used to route NFS traffic between compute nodes and NFS file servers on the 1/10 Gb/s Ethernet infrastructure.

3.2.8.3 Gateway Node I/O Configuration (TR-1)

The SU gateway nodes chipset will be configured with sufficient PCIe 8x (or faster) or HyperTransport™ busses and sufficient PCIe 8x (or faster) or HTX slots to drive the IBA HCA and one or two 10 Gb/s HCA with a total of two 10 Gb/s Ethernet ports that can deliver at least 600 MB/s each. The SU gateway node PCIe and HyperTransport™ infrastructure will be fully compatible with the proposed IBA HCA and 10 Gb/s Ethernet cards.

3.2.8.4 Gateway Node 10 Gb/s Ethernet Card (TR-1)

The SU gateway nodes will be configured with one (or two) PCIe 8x (or faster) dual (or single) port 10 Gb/s IEEE 802.3ae 10 GBase-SR with multi-mode fiber Ethernet card(s). Offeror will provide and support an open source 10 Gb/s Ethernet driver for Linux 2.6 kernels. The 10 Gb/s Ethernet and device driver will support 9KB jumbo frame operation. The 10Gb/s Ethernet card(s) will have additional capabilities such as: TOE, SDR, RNIC and iSER.

3.2.8.5 Gateway Node Delivered Performance (TR-2)

The SU gateway nodes will be configured to support a minimum of 1.2 GB/s. IBA to 10 Gb/s Ethernet delivered IP routing bi-directional bandwidth (counting both directions). Offeror will provide fully documented benchmark data demonstrating the minimum performance utilizing the NTTCP benchmark with the Offeror's response.

3.2.9 Login/Service/Master Node Requirements

The following Requirements are specific to the Login/Service/Master (LSM) nodes and supersede the general node requirements, above. LSM nodes will be used for management functions as well as user access (e.g., Login, application development and job launch).

3.2.9.1 LSM Node Count (TR-1)

The Offeror will configure the SU with one LSM node for SU peak between 10 teraFLOP/s and 20 teraFLOP/s and two LSM nodes for SU peak above 20 teraFLOP/s.

3.2.9.2 LSM Node I/O Configuration (TR-1)

The SU LSM nodes chipset will be configured with sufficient PCIe 8x (or faster) or HyperTransport™ busses and sufficient PCIe 8x or HTX slots to drive the IBA HCA, 1 and 10 Gb/s Ethernet cards at full line rate. The SU LSM node PCIe or HyperTransport™ infrastructure will be fully compatible with the proposed IBA HCA, 1 and 10 Gb/s Ethernet cards.

3.2.9.3 LSM Node Ethernet Configuration (TR-1)

The SU LSM nodes will be configured with one (1) PCIe 8x (or faster) 10 Gb/s IEEE 802.3ae 10 GBase-SR with multi-mode fiber Ethernet card for access to the site 10 Gb/s Ethernet infrastructure. The SU LSM nodes will be configured with dual 1 Gb/s IEEE 802.3z 1000Base-SX multi-mode fiber ports supported by a PCIe bus for access to the site 1 Gb/s Ethernet infrastructure. These 1 and 10 Gb/s Ethernet ports will be in addition to any ports required for management functions. Offeror will provide and support an open source 1 and 10 Gb/s Ethernet driver for Linux 2.6.9 and later kernels. The 1 and 10 Gb/s Ethernet and device driver will support 9KB jumbo frame operation.

3.2.9.4 LSM Node Accessory Configuration (TR-2)

The SU LSM nodes will be configured with one (1) read only (not read/write) 4x DVD-ROM bootable device. The SU LSM nodes will be configured with one (1) 750 GB (or larger) SATA disk. This disk should be hot swappable and directly accessible from the exterior of the LSM node.

3.2.9.5 LSM Node Delivered Performance (TR-2)

The SU LSM nodes will be configured to support a minimum of 1.0 GB/s IBA to dual 1 and one 10 Gb/s Ethernet delivered IP bi-directional routing bandwidth (counting both directions). Offeror will provide fully documented benchmark data demonstrating the minimum performance utilizing the NTTCP benchmark with the Offeror's response.

3.2.10 Remote Partition Service Node Requirements

The following Requirements are specific to the Remote Partition (RPS) nodes and supersede the general node requirements, above. RPS nodes will be used as a remote disk device for the compute and gateway.

3.2.10.1 RPS Node Count (TR-1)

The Offeror will configure the SU with one (1) RPS node for SU peak between 10 teraFLOP/s and 20 teraFLOP/s and two (2) RPS nodes for SU peak above 20 teraFLOP/s.

3.2.10.2 RPS Node I/O Configuration (TR-1)

The SU RPS nodes chipset will be configured with sufficient PCIe 8x (or faster) or HyperTransport™ busses and sufficient PCIe 8x or HTX (or faster) slots to drive the IBA HCAs, 1 Ethernet card at full line rate. The SU RPS node PCIe and HyperTransport™ infrastructure will be fully compatible with the proposed IBA HCA, 1 Ethernet cards.

3.2.10.3 RPS Node Ethernet Configuration (TR-1)

The SU RPS nodes will be configured with two (2) 1 Gb/s IEEE 802.3ab 1000BaseT copper ports supported by a PCIe bus for access to the management Ethernet infrastructure. These 1 Gb/s Ethernet ports will be in addition to any ports required for RPS node management functions. Offeror will provide and support an open source 1 Ethernet driver for Linux 2.6.9 and later kernels. The 1 Ethernet and device driver will support 9KB jumbo frame operation.

3.2.10.4 RPS Node RAID5 Configuration (TR-1)

The SU RPS nodes will be configured with at least one (1) highly reliable hardware RAID5 in (D+P) configuration utilizing 750 GB SATA (or better) disks. The number of data drives, D will be at least 7. The RAID5 will be configured with one (1) parity (P) drive. The capacity of these RAID5 arrays, counting only the D drives) will be at least (M+2)*(C+G) GB, where M is the memory size of the compute and gateway nodes, C is the number of compute nodes and G is the number of gateway nodes in the SU. The RAID5 arrays will deliver at least 500 MB/s large block read/write bandwidth from the Linux EXT3 filesystem mounted on each partition. The RAID5 arrays will deliver at least 640 IOPS to an IO workload randomly reading and writing 4,096B block from the Linux EXT3 filesystem mounted on each partition.

The proposed RAID5 arrays will be compatible with the Linux 2.6 LVM and paritionable into 2*(C+G) partitions, half of the partitions with size M GB and half of the partitions of size 2 GB with LVM. The RAID5 device will integrate into the node management infrastructure and be capable of notifying system administrators when a drive has failed. The RAID5 array will be configurable to automatically rebuild the RAID5 array using the hot spare drive. This disk should be hot swappable and directly accessible from the exterior of the RPS node.

Example: A 20 teraFLOP/s SU with 144 nodes total, 1 LSM node, 1 RPS node and 6 gateway nodes will have 136 compute nodes. Suppose the four socket nodes in the SU have a peak of 140.8 gigaFLOP/s and 32 GB of memory. Then the RAID5 capacity of the RPS node will be at least (32+2)*(136+6) = 4,828 GB. The capacity requirement will be satisfied with 7x750 GB SATA data drives.

3.2.11 SU Management Ethernet (TR-1)

The Offeror will provide a management Ethernet (100 BaseT or faster, 1000Base-TX (copper) is preferred) for the SU. The management Ethernet infrastructure will provide access to every node. In the case of failure in the IBA interconnect, the management network can be used to boot the entire system. The management Ethernet will be aggregated with high quality, high reliability Ethernet switches with full bandwidth backplanes and provide a single 1000Base-TX (copper) Ethernet uplink. The management Ethernet cables will be bundled within the rack in such a way as to not kink the cables, nor place strain on the Ethernet connectors. All management Ethernet connectors will have a snug fit when inserted in the management Ethernet port on the nodes and switches. The management Ethernet cables will meet or exceed Cat 5E specifications for cable and connectors. Cable quality references can be found at: (http://www.integrityscs.com/index.htm) and (http://www.panduitncg.com:80/whatsnew/integrity_white_paper.asp).
If 1000Base-TX (copper) is offered, then Offeror will provide CAT6 or equivalent management cables. A suggested source of this quality cable is Panduit corporations Powersum+ tangle free patch cords, Part# UTPCI10BL for a 10' cable. The URL for this product is: (http://www.panduitncg.com/solutions/copper_category_5e_5_3.asp).
Management Ethernet reliability is specified in Section 4.1.

3.2.12 SU Racks and Packaging (TR-1)

The Offeror will place the TLCC SU nodes, global disks, RAID controllers, IBA infrastructure, management Ethernet, in standard computer racks with ample room for cable management of InfiniBand 4x cables, CAT5e or CAT6 Ethernet cables and console serial cables and power cables. There will not be any console display, keyboard or mouse (KVM) equipment in the racks, except in the rack containing the LSM node. The LSM node, in each SU, will be connected to a single rackmount 1U keyboard, monitor and mouse.

3.2.12.1 SU Design Optimization (TR-2)

Offeror's SU design will be optimized to minimize the overall footprint of 2, 4, 6 and 8 SU aggregations within the other constrains in section 3.2.12.

3.2.12.2 Rack Height and Weight (TR-1)

SU will not be taller than 84" high (48U) and not place an average weight load of more than 250 lbs/ft^2 over the entire footprint of the SU, including hot and cold isles. If Offeror proposes a rack configuration that weighs more 250 lbs/ft^2 over the footprint of the rack, then Offeror will indicate how this weight can be redistributed over more area to less than 250 lbs/ft^2.

Comment [MKS1]: The following LANL requirement needs to be rewritten so that it is intellagable. Also, this requirement is specific to LANL.
Rack Cable Signal Integrity (TR-1)
All cables in the SU shall undergo testing that abides by the EIA/TIA 568 B standards.

The testing of fiber shall consist of determining the db loss of the cable at the 850 nm and 1300 nm wavelengths for multi-mode fiber, and 1310 nm and 1550 nm wavelengths for the single mode fiber. Each test shall be bi-directional (testing from each end). The specifications for the maximum loss depends on the number of connectors tested and the length of the cable. The maximum db loss per connector is 0.75 db, while for fiber loss per kilometer (length of cable) it is 3.5 db at 850 nm, 1.5 db at 1300 nm, and 1.0db for 1310 nm and 1550 nm. The MPO fiber InfiniBand cable shall also meet these requirements. Each fiber cable shall have a unique serial number and label (see Section 4.4) that is mapped to the above fiber test results.

For the testing of copper cable, a wire map test shall be performed to make sure the individual conductors are not crossed, a test to determine the cross talk measurements (NEXT), return loss measurements and signal to noise comparisons. Each copper cable shall have a unique serial number and label (see Section 4.4) that is mapped to the above fiber test results. Since there is no feasible way to test InfiniBand 4X copper cable today, it is currently exempt from these requirements.

There is testing equipment that have all these parameters either pre-set or the capability of being set for a particular type of cable (copper or fiber) that already meet the EIA/TIA standards. The most common are power meters or Optical Time Domain Reflectometers (OTDR).

For

3.2.12.3 Rack Structural Integrity (TR-2)

The provided racks will be of high structural quality. In particular, rack frames will be of sufficient strength and rigidity that the racks will not flex nor twist under the external load of a human being pushing at eye level on the rack from any of the four corners or sides. Additional reinforcement will be added as necessary to maintain rack structural integrity.

As a seismic event precaution, upon SU delivery Offeror will bolt racks in each row together with at least four 3/4" lag bolts, or better, for end racks and eight 3/4" lag bolts, or better, for racks sited touching two other racks (one bolt on each side corner touching another rack). During SU assembly at Offeror's facility, racks should have holes for inter rack bolting drilled prior to the placement of ANY equipment in them. SU's sited next to existing equipment (e.g., prior SU deliveries) in the same row need not be bolted together. The rack base will have wheels, leveling feet and adequate structural integrity to allow the rack to be bolted internally through the computer floor to the concrete sub-floor where required. The rack base must also allow adequate hole penetration for power and communication cables.

3.2.12.4 Rack Air Flow and Cooling (TR-1)

The racks will have sufficient air flow to adequately cool at full load the equipment mounted in the rack and racks installed at 600 ft. (LLNL) and 7,500 ft. (LANL or SNL) elevation with 30% humidity at up to 60º F (LLNL) and 80º F (LANL or SNL) air intake temperature. The rack bottom panel will be completely removed to improve airflow and allow sufficient room to run IBA cabling out of the rack and under the floor.

3.2.12.5 Rack Doors (TR-2)

The rack, if provided with a front or rear door, will include a clear or tinted non-breakable plexi-glass (or similar material) panel and have sufficient perforations to maintain adequate air flow throughout the cabinet while closed. The front and rear rack doors will be lockable.

3.2.12.6 Rack Cable Management (TR-2)

The racks will have sufficient room for all equipment and cables without impeding airflow through the rack. All cables within a rack will be supported so that the weight of the cable is not borne by the cable attachment mechanism. A rubber grommet or other protection will be placed around the rack bottom opening as necessary to protect the IBA and other cables from damage. In addition, cables will be attached to rack mounts installed in the rear and/or front of the cabinet for cable management. Cable management solution may not block access to active components in the rack. Rack cabling may allow the removal of any FRU in the rack without having to significantly uncable or recable the entire rack.

3.2.12.7 Rack Color (TR-3)

All racks will be black and covered with a fully powder coated style paint or other covering.

3.2.12.8 Rack Power and Cooling (TR-1)

Overall power and cooling for the SU are TCO components for the Tri-Laboratory

community. For racks with air cooling solutions that require all the cooling from air provided by the facility, each rack will not require more than 50 kW (LLNL), 20 kW (LANL or SNL) of power, and corresponding cooling, assumin g front to back air cooling. If the rack requires more than the above power envelopes, then Offeror will propose less dense solutions and/or alternative cooling apparatus that reduces the intake air cooling load. Offeror will fully describe the liquid cooling apparatus and the implications for siting and facilities modifications (e.g., chilled water feeds, flow rates, etc).

3.2.12.9 Rack PDU (TR-1)

Rack PDU for the SU will minimize the number of 208V circuit breakers required in wall panels at the Tri-Laboratory sites. One (1) would be ideal. In addition, the amperage of the required circuit breakers should be calibrated so that the utilization is maximized, but below 80% of the rated load during normal operation with user applications running. If the equipment in the rack requires more power during power-up (so called surge power), the rack PDU may not be calibrated to this surge power, but rather to the normal operating power with user applications running.

The Rack PDU will have on-off switches or switch rated circuit breakers to allow system administrator to power down all components in a rack with switches or circuit breakers in the PDU.

3.2.13 UL Listing (TR-1)

All supplied electrical components for the SU will be UL-Listed or certified by the Canadian Standards Association (or recognized equivalent). In addition, all rack internal wiring will be designed and implemented to National and State (CA, NM) electrical codes.

3.3 TLCC Software Requirements (TR-1)

This section describes the software requirements beyond the Government Furnished Software (GFS) for the TLCC. The software associated with building and installing the TLCC SUs is described in Sections 2.3. Offeror will provide all source code as Open Source in the form of buildable source SRPMs with the provided software.

3.3.1 Minimum IBA Software Stack (MR)

The Offeror shall provide and support a fully compliant IBA V1.2 (http://www.infinibandta.org/specs) Linux software stack for the TLCC SU. The Offeror's IBA software stack shall be fully functional, stable and scale on clusters comprised of aggregations of 1 to 8 SUs. The IBA software stack shall include at least the following components: HCA driver, core InfiniBand modules (subnet management agent, performance management agent, connection manager, subnet administration client, general service interface, queue-pair redirection), kernel and user space Verbs, IPoIB, SRP or iSER (initiator and target), MPI, OpenSM, network and host diagnostics and firmware for HCA and IBA switches and Linux command line utilities for flashing the HCA firmware from the node it is attached to and the switch firmware LSM node over the management Ethernet.

3.3.1.1 OpenFabrics Enterprise Distribution (MR)

The Offeror's supplied and supported IBA software stack shall be OpenFabrics

Enterprise Distribution (OFED) version 1.2 or later. The supplied and supported OpenFabrics software stack shall be certified as OpenFabrics compliant after successfully passing the OpenFabrics compliance test suite and being released by the OpenFabrics Alliance. http://www.openib.org/docs/openib_compilance_test_suite.html The Offeror's OFED software stack shall be deemed production quality by the Tri-Laboratory community if it successfully completes the Tri-Laboratory (pre-ship, post-ship and/or acceptance) workload test plan exit criteria on the proposed hardware at 1 SU scale.

The Offeror shall contribute all modifications to the OFED software stack to the OpenFabrics Alliance throughout the lifetime of this procurement. The Offeror shall document and track all their OFED software stack bugs in the OpenFabrics Alliance bugzilla bug tracking system (http://openfabricshttp://openib.org/bugzilla.org/bugzilla).

3.3.2 IBA Upper Layer Protocols (TR-1)

The Offeror's provided and supported OFED stack releases will also include the following Upper Layer Protocols (ULP):
SDP (www.rdmaconsortium.org/home)
user space DAPL, http://www.datcollaborative.org/udapl.html
IPoIB, http://www.ieft.org/html.charters/OLD/ipoib-charter.htmlhttp://www.datcollaborative.org/kdapl.html
SRP, http://www.t11.org/t10/drafts/srp/srp-r16a.pdf
iSCSI, http://www.ietf.org/rfc/rfc3720.txt
iSER, http://www.rdmaconsortium.org/home
NFS-RDMA, http://www.ietf.org/rfc/rfc3010.txt
IPoIB connected mode, http://www.ietf.org/internet-drafts/draft-ietf-ipoib-connected-mode-00.txt

These protocols will fully implement and conform to the above specifications. The Offeror's OFED ULPs will successfully pass all relevant tests in the OpenFabrics compliance test suite.

3.3.3 TLCC Peripheral Device Drivers (TR-1)

Offeror will provide Linux drivers for all peripheral devices supplied that function with the LLNL, LANL and SNL kernels. Offeror will specifically call out and fully disclose any proposed peripheral device drivers required with the proposed SU including version number and provide system administration or programmers documentation with the proposal.

3.3.4 RPS Node Software (TR-1)

The RPS node will provide remote access to root and /swap file systems for compute and gateway nodes. As the compute and gateway nodes boot over the management or IBA network, they will establish connections to the RPS node and mount their root (including /tmp and /var/tmp) (EXT3 file system) and /swap (block device) partitions from the RPS node via SCSI Remote Protocol (SRP) or iSCSI with extensions for RDMA (iSER) target over IBA. Thus, the Offeror will supply the SRP or iSER target code for the RPS node. In addition, Offeror will supply the RAID5 device driver for the RPS node RAID5 controller that is compatible with the RHEL4 LVM layer. With this RPS node software configuration

it will be possible to simultaneously service root and /swap partitions for all compute and gateway nodes.

3.3.5 Kernel Support for Memory Error Detection (TR-1)

The Offeror will modify the Error Detection and Correction (EDAC) code in the Linux kernel or loadable kernel module to support the chipsets proposed (Section 3.2.5.7). See http://sourceforge.net/projects/bluesmoke/ for more information on EDAC. The Offeror will work with the Tri-Laboratory community to integrate this code into the BASS and Tri-Laboratory Linux distributions. This EDAC software with Offeror supplied modifications will log all correctable and uncorrectable memory errors to the Linux kernel log facility. This modification will report the failed or failing DIMM FRU (i.e., the exact DIMM location on the motherboard will be indicated in the kernel panic message). This EDAC software with Offeror supplied modifications will panic the node if the memory subsystem generates an uncorrectable memory error. This EDAC software with Offeror supplied modifications will provide an appropriate interface to the hardware diagnostics in Section 4.5.2.

3.3.6 Linux Access to Motherboard Sensors (TR-1)

If an IPMI 2.0 and BMC remote management solution (section 3.2.7.1) is proposed all IPMI sensor data will be accessible in-band and out-of-band through FreeIPMI. Offeror will provide any changes required to FreeIPMI. These changes will be provided by the Offeror for inclusion in the open source FreeIPMI project.

If a TRMS (section 3.2.7.1) is proposed, Offeror will provide at least a command line mechanism for sampling motherboard sensor values for those listed in section 3.2.7.1 from within the Linux operating system. Sensor types and values will be output in such a way that they can be parsed by scripts. The LM-SENSORS package (http://secure.netroedge.com/~lm78/) is one solution used by the Tri-Laboratory community. If LM-SENSORS is proposed, Offeror will provide any needed kernel device drivers under open source license and a correctly calibrated sensors.conf file, including threshold values that adhere to manufacturers specifications, for all node types offered.

3.3.7 Remote Management Software (TR-2)

The Offeror will provide remote management software, beyond that defined in Section 3.2.7, for the remote management of the TLCC Cluster. This may include utilities to capture and monitor BIOS, Linux Console and other node management I/O. It is preferred that any provided software be Open Source.

3.3.7.1 TRMS Software (TR-1)

Offeror will supply software that handles remote power control over LAN and interfaces with PowerMan software. Offeror will deliver software that handles remote console access and logging over LAN and interfaces with the ConMan software.

3.3.7.2 IPMI and BMC Remote Management Software (TR-1)

If the Offeror proposes IPMI and BMC remote management solution (section 3.2.7.2), then Offeror will supply FreeIPMI software as the software component of the IPMI and BMC remote management solution (section 3.2.7.2), if proposed.

3.3.7.3 Linux Tool for LinuxBIOS Setup (TR-1)

The node BIOS will be delivered with a Linux command line tool to save/restore BIOS setup parameters. This tool will allow individual BIOS parameters to be modified. This tool will be provided to the Tri-Laboratory community under open source license. It is not acceptable to deliver a BIOS setup tool that only works in an operating system other than Linux (e.g., DOS or Windows).

3.3.7.4 Linux Tool for BIOS Upgrade (TR-1)

The node BIOS will be delivered with a Linux command line tool for BIOS upgrade.

If Offeror proposes a standard BIOS, the MTD kernel device driver (http://www.linux-mtd.infradead.org) is one solution to this requirement that is already used by the University. MTD offers a UNIX character device interface to common flash memory technology devices. On the TLCC SUs, an MTD based solution would be combined with scripts that safely implement flash/verify functions for the Offeror's BIOS images. If this method is proposed, all modifications to the MTD device driver and scripts will be provided to the Tri-Laboratory community under open source license. If another mechanism to meet this requirement is proposed, then the proposed tool will be offered to the Tri-Laboratory under open source license. It is not acceptable to deliver a tool that only works in an operating system other than Linux (e.g., DOS or Windows).

3.3.8 System Diagnostics (TR-2)

See Section 4 for the list of system monitoring and diagnostics required.

3.4 SU Hardware Evolution (TR-1)

For SU technology components that Offeror rated with "medium" or "high" impact in section 3.1.3, Offeror will list how those changes to the proposed solution will change the offering relative to the hardware requirements in section 3.2. Offeror's response should indicate changes to meeting each requirement in section 3.2.X, with designation 3.4.X, below.

3.5 SU Software Evolution (TR-1)

For SU technology components that Offeror rated with "medium" or "high" impact in section 3.1.3, Offeror will list how those changes to the proposed solution will change the offering relative to the software requirements in section 3.3. In addition, any software component offering that changes over the 3QCY07 through 3QCY08 time frame that has "medium" or "high" impact should be listed as well. Offeror's response should indicate changes to meeting each requirement in section 3.3.X, with section heading 3.5.X, below.

End of Section 3

4 Reliability, Availability, Serviceability (RAS) and Maintenance

The TLCC SUs, aggregated into clusters, are intended for classified production usage at the Tri-Laboratory sites. The Tri-Laboratory's therefore requires that the SUs have highly effective, scalable RAS features and prompt hardware and software maintenance. In addition, Offeror shall propose end-to-end support for the proposed IBA interconnect hardware and software.

For hardware maintenance, the strategy is that Tri-Laboratory personnel will provide on-site, on-call 24×7 hardware failure response. We envision that these hardware technicians and system administrators will be trained by the Offeror to perform on-site service on the provided hardware. For easily diagnosable node problems, Tri-Laboratory personnel will perform repair actions in-situ by replacing FRUs. For harder to diagnose problems, Tri-Laboratory personnel will swap out the failing node(s) with on-site hot spare node(s) and perform diagnosis and repair actions in the separate Hot-Spare Cluster (HSC). Failing FRUs or nodes (except for writable media) will be returned to the Offeror for replacement. Hard Disks FRUs and writeable media (e.g., EEPROM) from other FRUs will be destroyed by each Laboratory according to DOE/NNSA computer security orders. Thus, the Tri-Laboratory requires an on-site parts cache of all FRUs and a small cluster of fully functional hot-spare nodes of each node type. The Offeror will work with the Tri-Laboratory community to diagnose hardware problems (either remotely or on-site, as appropriate). On occasions, when systematic problems with the cluster are found, the Offeror's personnel will augment Tri-Laboratory personnel in diagnosing the problem and performing repair actions.

In order for the large number TLCC SUs to fulfill the mission of providing the capacity resource for ASC and SSP, they must be highly stable and reliable from both a hardware and software perspective. The number of failing components per unit time (weekly) should be kept to a minimum. SU components should be fully tested and burned in before delivery (initially and as FRU or hot-spare node replacement). In addition, in order to minimize the impact of failing parts, the Tri-Laboratory community must have the ability to quickly diagnose problems and perform repair actions. A comprehensive set of diagnostics that are actually capable of exposing and diagnosing problems are required. It has been our experience that this is a difficult but achievable goal, and the Offeror will need to specifically apply sufficient resources to accomplish it.

For software, the strategy is similar to the hardware strategy in that Tri-Laboratory personnel will perform the Level 1 and Level 2 software support functions. Specifically, Tri-Laboratory personnel will diagnose software bugs to determine the failing component. The problem will be handed off to the appropriate Tri-Laboratory organization for resolution. For Tri-Laboratory supplied system tools, Tri-Laboratory personnel will fix the bugs. For Offeror-supplied system tools, the Offeror will need to supply problem resolution. For the Linux kernel and associated utilities, the Tri-Laboratory community intends to separately subcontract with Red Hat for Enterprise level support. For file system related SW problems, the Tri-Laboratory community intends to separately subcontract with Cluster File Systems, Inc for Lustre support and with Panasas for PanFS support. For compilers, debugger and application performance analysis tools, the Tri-Laboratory community intends to separately subcontract with the appropriate vendors for support.

4.1 Highly Reliable Management Network (TR-1)

The SU management Ethernet will be a highly reliable network that does not drop a single node from the network more than once a year. For example an SU design with 144 nodes, the connection between any TLCC SU node and the management network will be dropped less than once every 1,780 months. This is both a hardware and a software (Linux Ethernet device driver) requirement. In addition, the management network will be implemented with connectors on the node mating to the management Ethernet cabling and connectors (Section 3.2.11) so that manually tugging or touching the cable at a node or switch does not drop the Ethernet link. The management Ethernet switches (Section 3.2.11) will be configured such that they behave as standard multi-port bridges.

Each SU management Ethernet will be connected via one 1Gb/s Ethernet uplink

4.2 TLCC Node Reliability and Monitoring (TR-2)

The TLCC SU nodes and other hardware components will have a Mean Time Between Failure (MTBF) of greater than 193,536 hours (less than one node failure per week per 1,152 nodes). Any failing SU hardware component that causes one or more nodes to becoming unavailable for job scheduling or kills the job running on it is included in this MTBF statistic. For example failures of: DIMMs, processors, motherboards, non-redundant power supplies, blade chassis, IBA infrastructure, PDUs all cause nodes to crash or make nodes become unavailable and are counted in this statistic. Redundant parts that fail such as power supplies, fans, single memory chips that are covered by chip-kill, but do not cause nodes to crash or become unavailable do not count in this statistic. Parts removed under preventive maintenance prior to actual failure such as DIMMs removed after excessive correctable single bit errors or SATA drives removed due to large number of block remaps or SMART future failure indicators do not count in this statistic.

The TLCC SU nodes will have real-time hardware monitoring, at a specified interval, of system temperature, processor temperature, fan rotation rate, power supply voltages, etc. This node hardware monitoring facility will alert the scalable monitoring software in Section 4.6 via serial console or management network when any monitored hardware parameter falls outside of the specified nominal range. In addition, the system components may provide failure or diagnostic information via the serial console or management network.

4.3 In Place Node Service (TR-1)

The SU nodes will be serviceable from within the rack. The node will be mechanically designed to have minimal tool requirements for disassembly. The nodes and other rack components will be mechanically designed so that complete node other rack component disassembly and reassembly can be accomplished in less than 20 minutes by a trained technician without having to move the rack. Blade solutions will have hot swappable blades: the blade chassis will not require being powered down during a blade replacement repair action.

4.4 Component Labeling (TR-1)

Every rack, Ethernet switch, Ethernet cable, IBA switch, IBA cable, node, disk enclosure will be clearly labeled with a unique identifier visible from the front of the rack and/or the rear of the rack, as appropriate, when the rack door is open. These labels will be high quality

so that they do not fall off, fade, disintegrate, or otherwise become unusable or unreadable during the lifetime of the cluster. The font will be non-serif such as Arial or Courier with font size for these labels at least 9pt. 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. This serial number needs to be visible without having to disassemble the node, or else it will be queryable, either from Linux or the BIOS from a Linux command line tool.

4.5 Field Replaceable Unit (FRU) Diagnostics (TR-2)

Diagnostics will be provided that isolate a failure of a TLCC SU component to a single FRU for the supplied hardware. These diagnostics will run from the Linux command line. The diagnostic information will be accessible to operators through networked workstations.

4.5.1 Node Diagnostics Suite (TR-1)

The Offeror will provide a set of hardware diagnostic programs (a diagnostic suite or diagnostics) for each type of node provided that run from the Linux command line and produce output to STDERR or STDOUT and exit with an appropriate error code when errors are detected. These diagnostics will be capable of stressing the node motherboard components such as processors, chip set, memory hierarchy, on-board networking (e.g., management network), peripheral buses, and local disk drives. The diagnostics will stress the memory hierarchy to generate single and double bit memory errors. The diagnostics will read the hardware single and double bit memory error counters and reset the counts to zero. The diagnostics will stress the local disk in a nondestructive test to generate correctable and uncorrectable read and/or write errors. The diagnostics will read the hardware and/or Linux recoverable I/O error counters and reset the counts to zero. The diagnostics will stress the integer and floating point units in specific core(s) in serial (i.e., one processor and/or HyperThread, as appropriate, at a time) or in parallel (i.e., multiple processors and/or multiple HyperThreads, as appropriate). The CPU stress tests will bind to a specific core and/or HyperThread, as appropriate, by command line option, if possible.

4.5.2 Memory Diagnostics (TR-1)

The Linux OS will interface to the SU node hardware memory error facility specified in Section 3.3.5 to log all correctable and uncorrectable memory errors on each memory FRU. When the node experiences an uncorrectable memory error, the Linux kernel will report the failing memory FRU and panic the node. The Offeror will provide a Linux utility that can scan the nodes and report correctable and uncorrectable memory errors at the FRU level indicating the exact DIMM location on the motherboard where the failing or failed DIMM is located and reset the counters.

4.5.3 IBA Diagnostics (TR-1)

The Offeror will provide a set of IBA hardware diagnostic programs (a diagnostic suite or diagnostics) for IBA components provided that run from the Linux command line and produce output to STDERR or STDOUT and exit with an appropriate error code when errors are detected. These diagnostics will be able to diagnose failures with IBA HCA, cables, and switch hardware down to individual FRU. These diagnostics will run and correctly diagnose failed and intermittently failing hardware within four hours and find **all** failed or intermittently failing components. In addition, these IBA diagnostics will be able

detect slow portions or portions with high bad packet rates of the interconnect infrastructure.

4.5.4 Peripheral Component Diagnostics (TR-2)

The Offeror will provide a set of hardware diagnostic programs (a diagnostic suite or diagnostics) for each type of peripheral component provided that run from the Linux command line and produce output to STDERR or STDOUT and exit with an appropriate error code when errors are detected. At a minimum, the diagnostics will test the 1000BaseT and other provided networking (e.g., Fibre Channel) adapters, RAID device and disks.

4.6 Scalable System Monitoring (TR-2)

There will be a scalable system monitoring capability for the TLCC SU supplied by the Tri-Laboratory Community that has a command line interface (CLI) for scriptable control and monitoring. This facility will directly interface to the Offeror provided node monitoring facility through the ”Motherboard Sensors” Section 3.3.6 facilities. All SU monitoring information will be centrally collected by the Tri-Laboratory Community at intervals set by the system administrator on one of the service nodes. All SU monitoring and diagnostics information provided by the Offeror will be formatted so that “expect scripts” can detect failures. In addition, this facility will be used by the Tri-Laboratory Community to launch diagnostics in parallel over the management network across all or part of the cluster, as directed by a system administrator from the Linux command line on one of the service nodes.

4.7 Hardware Maintenance (TR-1)

The Offeror will supply hardware maintenance for each proposed TLCC SU for a three-year period starting with cluster acceptance. Note that this implies that the number of SU under maintenance will ramp up over the delivery schedule and ramp down starting three years after the first TLCC07 cluster is accepted. Tri-Laboratory personnel will attempt on-site first-level hardware fault diagnosis and repair actions. Offeror will provide second-level hardware fault diagnosis and fault determination during normal business hours. That is, if Tri-Laboratory personnel cannot repair failing components based on-site parts cache, then the Offeror personnel will be required to make on-site repairs. Offeror supplied hardware maintenance response time will be before the end of the next business day from incident report until Offeror personnel perform diagnosis and/or repair work. The proposed system will be installed in a limited access area vault type rooms (VTR) at the Tri-Laboratory sites and maintenance personnel must obtain DOE P clearances for repair actions at LLNL and be escorted during repair actions. USA Citizenship for maintenance personnel is highly preferred because it takes at least 30 days to obtain VTR access for foreign nationals from non-sensitive countries. During the period from the start of SU installation through acceptance, Offeror support for hardware will be 12 hour a day, seven days a week (0800-2000 PDT for LLNL and 0800-2000 MDT for LANL and Sandia), with one hour response time.

4.7.1 On-site Parts Cache (TR-1)

A scalable parts cache (of FRUs and hot spare nodes of each type proposed) at Tri-Laboratory sites is required that will be sufficient to sustain necessary repair actions on all proposed hardware and keep them in fully operational status for at least one week without parts cache refresh. That is, the parts cache, based on Offeror’s MTBF estimates for each FRU and each SU, will be sufficient to perform all required repair actions for one week

without the need for parts replacement and should be scaled up as SUs are delivered. The Offeror will propose sufficient quantities of FRUs and hot-spare nodes for the parts cache. The parts cache will be enlarged, at the Offeror's expense, should the on-site parts cache prove in actual experience to be insufficient to sustain the actually observed FRU or node failure rates. However, at a minimum, the on-site parts cache will include the following fully configured (except for IBA HCA) nodes: ten compute nodes and one LSM node. The Offeror will supply sufficient racks and associated hardware/software to make the HSC a cluster that can run diagnostics on every HSC node over the management Ethernet. In addition, a minimum of the following parts (and quantity), if bid: SATA Disks (2), SDRAM DIMM kit for a node (5), power supplies of each type (10), fans of each type (10), management Ethernet switch (1) and TRMS FRUs (1). The Tri-Laboratory Community will administer the nodes as a separate HSC in the unclassified environment. The Tri-Laboratory Community will store and inventory the HSC and other on-site parts cache components. Parts in the parts cache are Tri-Laboratory property. Failed parts become Offeror's property when RMAed back to Offeror.

4.8 Software Support (TR-1)

The Offeror will supply software maintenance for each Offeror supplied software component, specifically including the supplied BIOS and/or LinuxBIOS, starting with the first SU acceptance and ending three years after the final SU acceptance. Offeror provided software maintenance will include an electronic trouble reporting and tracking mechanism and periodic software updates. In addition, the Offeror will provide software fixes to reported bugs. The electronic trouble reporting and tracking mechanism will allow the Tri-Laboratory community to report bugs and status bug reports 24 hours a day, seven days a week. The Tri-Laboratory community will prioritize software defects so that the Offeror can apply the software maintenance resources to the most important problems. During the period from the start of SU installation through acceptance, Offeror support for supplied software will be 12 hour a day, seven days a week (0800-2000 PDT for LLNL and 0800-2000 MDT for LANL and Sandia), with one hour response time.

4.9 Mean Time Between Failure (MTBF) Calculation (TR-1)

The Offeror will provide the MTBF calculation for each FRU and node type. The Offeror will use these statistics to calculate the MTBF for the provided aggregate TLCC SU hardware. This calculation will be performed using a recognized standard. Examples of such standards are Military Standard (Mil Std) 756, Reliability Modeling and Prediction, which can be found in Military Handbook 217F, and the Sum of Parts Method outlined in Bellcore Technical Reference Manual 332. In the absence of relevant technical information in the proposal, the Tri-Laboratory community will be forced to make pessimistic reliability, availability, and serviceability assumptions in evaluating the proposal.

End of Section 4

5 Facilities Information

A portion of a existing facilities at the Tri-Laboratory sites will be used for siting the TLCC SU the following sections give detailed facilities information at the Tri-Laboratory sites and facilities based requirements for SU and cluster aggregations of SU.

5.1 LLNL Facilities Information

At LLNL, SU will be sited in on both the east and west main computer floor in B453 (see Figure 1.4-1-3). This entire facility has approximately 18,000 ft^2 and 2.5 MW of power and associated cooling available for this purpose for the SU compute, networking and peripherals. The computer floor is 48" raised floor with 250 lbs/ft^2 loading capability. However, racks with up to 500 lbs/ft^2 floor loading can be accommodated with additional floor bracing. The overall SU aggregations average floor loading (including isles between rows) can not exceed 250 lbs/ft^2. In addition, rolling weight of racks during installation can not exceed 250 lbs/ft^2. Power will be provided to racks by under floor electrical outlets supplied by the University to Offeror's specifications. Circuit breakers and PDUs are available in wall panels that can be modified to Offeror's specifications. All other cables must be contained in cable trays supplied by the University to Offeror's specifications. Straight point-to-point cable runs can **NOT** be assumed. The University will provide floor tile cut to Offeror's specifications. In addition, it is anticipated that the Offeror's equipment will be placed in adjacent rows so that air intakes in racks from adjacent rows are abutting with Offeror's specified separations and hot air exhausts in racks from adjacent rack rows are abutting with Offeror's specified separations. That is, the racks will be placed so that there are HOT and COLD aisle ways between racks with chilled air entering in the COLD isles and warmed air exiting in the HOT aisles. Offeror's will describe any unique cooling solutions that allow for more efficient utilization of computer floor space and provide information on facilities impacts. The University will provide 2'x2' grated floor tiles with 80% void in "cold isles" to product up to 2,500 CFM airflow per tile.

During installation, racks will transit from Offeror's delivery trucks through 12' (W) x 12' (H) roller doors to an interior delivery dock. Note that the delivery dock height is 45" and can only accommodate one tractor-semitrailer rig at a time. The racks will transit down a 300' long 9' 8" (W) x 8' 6" (H) hallway. The racks must transit several doors of size 7' 10" (W) x 7' 10" (H) and ride a freight elevator up one floor. The freight elevator doors are 8' 4" (W) x 8'(H), the elevator area is 8' (W) x 12' (D) and the maximum loading of the freight elevator is 10,000 lbs. Racks may be staged on the B453 computer floor for unloading from packaging or unpackaged on the interior delivery dock.

This facility will need modifications for siting TLCC SU. SU siting cost are considered part of the TCO. The University anticipates the construction work to be completed within sixty days of final facilities configuration agreement with Offeror. The siting assumes the following TLCC layout.

5.2 LANL Facilities Information

At LANL, most SUs will be sited in the SCC, Building 2327. This entire facility has approximately 20,000 ft^2 and 3.2 MW of power and associated cooling available for this purpose for the SU compute, networking and peripherals. The computer floor is 42" raised floor with 300 lbs/ft^2 loading capability. There is a 16 ft. ceiling and an 18'6" ceiling plenum.

The overall SU aggregations average floor loading (including aisles between rows) can not exceed 250 lbs/ft^2. In addition, rolling weight of racks during installation can not exceed 250 lbs/ft^2. Power will be provided to racks by under floor electrical outlets supplied by the University to Offeror's specifications. All other cables must be contained in cable trays supplied by the University to Offeror's specifications. Straight point-to-point cable runs can **NOT** be assumed. The University will provide floor tile cut to Offeror's specifications. In addition, it is anticipated that the Offeror's equipment will be placed in adjacent rows so that air intakes in racks from adjacent rows are abutting with Offeror's specified separations and hot air exhausts in racks from adjacent rack rows are abutting with Offeror's specified separations. That is, the racks will be placed so that there are HOT and COLD aisle ways between racks with chilled air entering in the COLD aisles and warmed air exiting in the HOT aisles. The University will provide 2'x2' grated floor tiles with 80% void in "cold aisles" to product up to 806,400 CFM airflow total for the entire room. The average airflow per floor tile in the SCC is 800-1200 CFM.

A few SUs delivered to LANL will be sited in the LDCC, Building 1498. This facility has 24" raised computer floor with 250 lb/ft^2 loading capability. There is a 10 ft. ceiling. The average airflow per floor tile in the LDCC is 600-800 CFM.

Electrical power distribution in both buildings uses 208V/120V 3-phase PDUs.

5.3 SNL Facilities Information

At SNL, SU can be sited in any of the following facilities: buildings/room 880/X-50, 880/220, and 725/105 in New Mexico, and 912/097 in California. The 880/X-50 facility has approximately 1,900 ft^2 of floor space, 1,050 kW of power, and 360 tons of cooling. The computer floor is 18" raised floor with 250lb/ft^2 loading capacity. There is a 10 ft ceiling and 5 ft ceiling plenum. The 880/220 facility has 800 ft^2, 400 kW of power, and 130 tons of cooling. The computer floor and ceiling measurements are the same as the measurements for the 880/X-50 facility. The 725/105 facility has 4,000 ft^2 of floor space, 1,400 kW of power, and 400 tons of cooling. The computer floor is 3 ft raised floor with 250lb/ft^2 loading capacity. There is a 20 ft ceiling. The 912/097 facility has 1,750 ft^2, 800 kW of power and 250 tons of cooling. The computer floor is 18-24" raised floor with 250lb/ft^2 loading capacity. There is a 10 ft ceiling and 2-4 ft ceiling plenum.

The overall SU aggregations average floor loading (including aisles between rows) can not exceed 250lb/ft^2. Power will be provided by under floor electrical outlets supplied by the Offeror's specifications. Sandia will provide floor tiles cut to the Offeror's specifications. Sandia will provide 2'x2' grated floor tiles with 55% void for the "cold aisles".

5.4 Power Requirements (TR-1)

Power requirements will be fully disclosed by Offeror at the time of proposal submission. This information will be communicated in the Offeror's proposal. Offeror will provide for each rack type: the number of kW or kVA required, the number and type of power connections required, and anticipated electrical load. This information will be verified by joint written (e-mail and text files) and telephone conferences between the Offeror and the Tri-Laboratory community no more than two months prior to the first SU delivery.

5.5 Cooling Requirements (TR-1)

Cooling requirements will be fully disclosed by Offeror in the proposal. Offeror will provide for each rack type (the number of Btu or tons AC required) and any environmental requirements, such as temperature and/or humidity range requirements. This information will be verified by joint written (e-mail and text files) and telephone conferences between the Offeror and the Tri-Laboratory community no more than two months prior to the first SU delivery.

5.6 Floor Space Requirements (TR-1)

Floor space requirements for a cluster with an aggregation of four SU (with associated IBA spine switches) will be fully disclosed by Offeror at the time of proposal submission. This information will be communicated in the Offeror's proposal. Offeror will provide a detailed floor plan (system layout) diagram indicating rack placement and location of required electrical outlets. This information will be verified by joint written (e-mail and text files) and telephone conferences between the Offeror and the Tri-Laboratory community no more than two months prior to the first SU delivery.

5.7 Delivery Requirements (TR-1)

If Offeror has any delivery requirements these will be communicated to the Tri-Laboratory community in the proposal. TLCC SU will be physically located inside a Limited Access Area in a Vault Type Room (VTR). The Tri-Laboratory community will only provide access to the room to authorized personnel under Authorized Escort. All on-site personnel will be required to submit applications for access and be approved by standard University procedures prior to entry into this facility. All on-site personnel at LLNL will require being DOE P-cleared or P-clearable. RFP responses should indicate if the on-site team has members that are other than US Citizens. Physical access to this facility by foreign nationals from sensitive countries (www.llnl.gov/expcon/sensitive.html) will not be allowed. Dialup capability and Internet access to the system will be allowed before the system goes classified, but not afterwards. Authorized individuals may be allowed remote access for running diagnostics and problem resolution only while the system remains unclassified.

Unless otherwise indicated in Offeror's proposal, installation crews will work up to an eight (8) hour day Monday through Friday, 8:00 a.m. to 5:00 p.m. Longer days, differing shift start/end times and/or weekend shifts can be accommodated by the Tri-Laboratory community at Offeror's request at least one week prior to delivery.

5.7.1 SU Installation Time (TR-1)

Offeror will deliver, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over each SU to the Tri-Laboratory community for acceptance testing within **three days** from the time the first truck delivering the SU backs up to the loading dock.

End of Section 5

6 Project Management

This procurement envisions purchasing a large number of Scalable Units during the duration of this project. The exact number of SUs delivered depends on the peak of the SU. In order to provide DOE/NNSA HQ the maximum amount of budget flexibility, additional SU are required beyond the base SU. The SU deliveries designated with "Option" in them are required as options that will be exercised at the sole discretion of the University upon request by the DOE/NNSA ASC HQ program office or individual Laboratory. With delivery of SUs to specific sites, additional 288 port, or larger, spine switches and cables are required as part of this RFP in order to aggregate SU into clusters. The receiving sites are responsible for providing 1-10 Gb/s Ethernet switching infrastructure to integrate the SU into the sites 1-10 Gb/s Ethernet backbones. The receiving sites are also responsible for integrating the delivered SU's into the sites existing multi-cluster file systems.

The construction, ship testing, delivery, installation, and acceptance testing of the TLCC07 SUs is a complex endeavor. It is anticipated that this project will require close coordination of Tri-Laboratory community, IBA supplier and the Offeror's personnel.

6.1 Open Source Development Partnership (TR-2)

The Offeror will provide information on the capabilities of the Offeror to engage in an Open Source development partnership and meet the goals set out in Section 1, 2, and 3 (i.e., OpenFabrics, Free IPMI, LinuxBIOS, OpenMPI). This information should include Offeror's financial health; Offeror's qualifications as a cluster provider; Offeror's qualifications as an Open Source development organization; cluster product roadmap and comparison to the overall TLCC strategy; the willingness of the Offeror to participate in the Open Source development, with other partners, of key missing HPTC cluster technology components such as scalable parallel file systems and cluster resource scheduling. If the Offeror has technology, such as a scalable parallel file system, cluster management tools, or cluster resource scheduling, that could be contributed to the Open Source community, please indicate that as well in the proposal.

6.2 Project Manager (TR-1)

The Offeror will provide the name and resume of the proposed project manager within the Offeror's corporation for the proposed activity. This project manager will be approved by the University technical representative. The project manager must be empowered by the Offeror's corporation to plan and execute the construction, shipment, and installation of the proposed configuration. This must include sufficient personnel and hardware resources within the corporation to assure successful completion of the activity on the proposed schedule.

6.3 Project Milestones (TR-3)

Delivery of the N SUs may be accomplished within four quarter year periods starting no later than 4QCY2007 and ending no later than 3QCY2008. Recall that N is defined in section 3.2.4. Let X = int(N/4) and Y = N-3*X.
Example: for N = 18 (20.3 teraFLOP/s SU), then X = 4 and Y = 6.

For the purposes of this section, the Tri-Laboratory community assumes a delivery schedule as follows:

Delivery Phase	Prototypes under separate contract	4QCY2007
1	Y SU	4QCY2007
2	X SU	1QCY2008
3	X SU	2QCY2008
4	X SU	3QCY2008
Total	N SU	

The Offeror may propose alternative SU delivery schedules within the start of the contract to the end of 3QCY2008.

The milestones below are structured so that a shared risk model with commodity price indexing is executed. Thus, Offeror and the Tri-Laboratory Community may review commodity pricing during the delivery period, as necessary. In addition, Offeror will deliver, additional spine switches, cables to enable the Tri-Laboratory personnel to combine multiple SU into larger clusters, as direction by the Tri-Laboratory in writing 30 days prior to SU delivery.

The Tri-Laboratory community plans to transition the TLCC SUs aggregated into clusters to classified operation Limited Availability (LA) Production status for a small group of ASC and SSP with fifteen days of delivery. The Tri-Laboratory community plans to run in LA status for a week and then transition the TLCC cluster to General Availability (GA) for the general community of ASC and SSP users. Instability of SU during acceptance, transition to LA or transition to GA will impede this schedule and should be avoided: time is of the essence.

In addition, there are multiple activities among multiple institutions and organizations within the institutions that must be coordinated prior to the first delivery and ongoing during the seven quarters of SU deliveries. In order to assure the timely execution of these programmatic goals and to make sure both parties understand the timeline, the Offeror will provide the Tri-Laboratory community with a project plan no more than seven days after subcontract award.

6.3.1 Detailed Project Plan (TR-1)

The Offeror will provide a detailed project plan no more than seven days after subcontract award. This project plan will include a Gantt chart with all the project milestones with dates and durations for work activities leading up to the milestones. The Gantt chart will indicate work activity and milestone and organizational dependencies. The Gantt chart will clearly indicate the project's critical path. At least one level of detail below each of the project milestones showing the work activities leading up to completion of the milestone will be included in the Gantt chart. The project plan will include a written Tri-Laboratory BASS build image checkout plan, pre-ship test plan and acceptance test plan. The project plan Gantt chart will be a Microsoft Project data file. The test plans will be Microsoft Word data files.

The build image checkout, pre-ship test and acceptance test plans will be mutually agreeable, but will include:
Successfully running with correct results three mixed MPI/OpenMP jobs (sPPM, UMT2K, LINPACK) sequentially or simultaneously across 90% of the SU compute nodes for at least four hours without failure;
Successfully running the LLNL Presta MPI stress test sequentially or simultaneously across 90% of the SU compute nodes for four hours without failure or performance anomalies; and
A demonstration that the Management Ethernet is functional, stable, and reliable.
Offeror will be responsible for LINPACK tuning and execution. The Tri-Laboratory community will be responsible for sPPM and UMT2K tuning and execution. The test plans will include clear test entry and exit criteria as well as a list of testing activities and benchmarks.

This milestone is complete when the University Procurement Representative (UPR) approves the project plan.

6.3.2 Tri-Laboratory BASS Final Checkout (TR-1, September 15, 2007)

The Tri-Laboratory community will provide the Build and Acceptance Software Stack (BASS) cluster distribution software (section 2.3.1) for installation on SU. After acceptance, the receiving Laboratory will install their local Linux distribution (see sections 2.3.2, 2.3.3 and 2.3.4) prior to moving the SUs into classified operation. The Offeror will assist the Tri-Laboratory community in development of provided components, test, debug and installation of these software stacks. This BASS and Laboratory local Linux Build Image effort will commence upon contract signing and continue throughout the contract time span. Offeror will assist the Tri-Laboratory community to finalize these software stacks prior to each SU manufacture. The BASS will then be used to manufacture, test, deliver and accept the SU. This milestone is complete when the first Offeror and Tri-Laboratory joint testing of the BASS Build Image completes the checkout test plan exit criteria by September 15, 2007.

6.3.3 TLCC07 Phase 1 Build (TR-1, October 2007)

The Offeror will build, fully assemble, configure, burn-in and test the Y SU for TLCC07 Phase 1 defined in section 6.3, as bid, with Tri-Laboratory Linux Build Images as directed by the UPR.

Offeror shall burn in and stress test TLCC07 Phase 1 equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 1 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC07 Phase 1 SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the University confirms that the correct BASS Build Image is installed on TLCC07 Phase 1 SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 1 SU meets the SWL pre-ship test entry criteria; 4) SWL pre-ship test is successfully executed on the TLCC07 Phase 1 SU; 5) The TLCC07 Phase 1 SU successfully completes the SWL pre-ship test exit criteria; 6) the UPR authorizes shipment of SU to tri-Laboratory sites; and 7) all equipment for this milestone leaves Offeror's integration location; 8) the required documentation is approved.

6.3.4 TLCC07 SU Phase 1 Delivery and Acceptance (TR-1, October 2007)

Offeror will deliver the TLCC07 SU to the Tri-Laboratory sites as directed by the UPR, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over the TLCC07 SU to the Tri-Laboratory community for acceptance testing. In addition, Offeror will deliver sufficient IBA spine switches and cables, as directed by the UPR, in order to allow the Tri-Laboratory community to assemble multiple SU clusters.

Offeror will burn in and stress test TLCC07 Phase 1 equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 1 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC07 Phase 1 SU with IBA hardware is installed at the tri-Laboratory sites, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the UPR confirms that the correct BASS Build Image is installed on the TLCC07 Phase 1 SUs; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 1 SUs meet the SWL post-ship test entry criteria; 4) SWL post-ship test is successfully executed on the TLCC07 Phase 1 SUs; 5) the TLCC07 Phase 1 SUs successfully completes the SWL post-ship test exit criteria; 6) the required documentation is approved.

6.3.5 TLCC07 Phase 1 Cluster Integration (TR-1, November 2007)

Offeror will deliver Phase 1 on-site hardware maintenance parts cache to each Phase 1 tri-Laboratory site. Offeror will integrate the site specific SU's into the single multi-SU (2, 4, or 8 SU) fully functional Clusters. The Offeror will fully assemble, pass University's delivery check list and initial functionality and performance verification testing, and turn over the integrated multi-SU Clusters to the University for integration testing.

Completion of this milestone starts the three year maintenance clock on the TLCC07 Phase 1 clusters.

This milestone is complete when: 1) all Cluster SU and IBA spine switch and cabling hardware is installed, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, the L1 and L2 IBA infrastructure must be functional); 2) the University confirms that the correct Laboratory software stack (see Sections 2.3.2, 2.3.3, and 2.3.4) is installed on the Cluster; 3) the UPR confirms that the Cluster on-site hardware maintenance parts cache is fully stocked; 4) the UPR confirms that the Cluster meets the integration test entry criteria; 5) the acceptance test is successfully executed on the Cluster; and 6) the Cluster successfully completes the integration test exit criteria; 7) the required documentation is approved..

6.3.6 TLCC07 Phase 1 Option Build (TR-1, October 2007)

At the Option of the University, the Offeror will build, fully assemble, configure, burn-in and test up to an additional 4 SU for TLCC07 Phase 1 Option, as bid, with Tri-Laboratory Linux Build Images as directed by the UPR.

Offeror shall burn in and stress test TLCC07 Phase 1 Option equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 1 Option cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC07 Phase 1 Option SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the University confirms that the correct BASS Build Image is installed on TLCC07 Phase 1 Option SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 1 Option SU meets the SWL pre-ship test entry criteria; 4) SWL pre-ship test is successfully executed on the TLCC07 Phase 1 Option SU; 5) The TLCC07 Phase 1 Option SU successfully completes the SWL pre-ship test exit criteria; 6) the UPR authorizes shipment of SU to tri-Laboratory sites; and 7) all equipment for this milestone leaves Offeror's integration location; 8) the required documentation is approved.

6.3.7 TLCC07 SU Phase 1 Option Delivery and Acceptance (TR-1, October 2007)

Offeror will deliver the TLCC07 SU to the Tri-Laboratory sites as directed by the UPR, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over the TLCC07 SU to the Tri-Laboratory community for acceptance testing. In addition, Offeror will deliver sufficient IBA spine

switches and cables, as directed by the UPR, in order to allow the Tri-Laboratory community to assemble multiple SU clusters.

Offeror will burn in and stress test TLCC07 Phase 1 Option equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 1 Option cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC07 Phase 1 Option SU with IBA hardware is installed at the tri-Laboratory sites, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the UPR confirms that the correct BASS Build Image is installed on the TLCC07 Phase 1 Option SUs; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 1 Option SUs meet the SWL post-ship test entry criteria; 4) SWL post-ship test is successfully executed on the TLCC07 Phase 1 Option SUs; 5) the TLCC07 Phase 1 Option SUs successfully completes the SWL post-ship test exit criteria; 6) the required documentation is approved.

6.3.8 TLCC07 Phase 1 Option Cluster Integration (TR-1, November 2007)

Offeror will deliver Phase 1 Option on-site hardware maintenance parts cache to each Phase 1 Option tri-Laboratory site. Offeror will integrate the site specific SU's into the single multi-SU (2, 4, or 8 SU) fully functional Clusters. The Offeror will fully assemble, pass University's delivery check list and initial functionality and performance verification testing, and turn over the integrated multi-SU Clusters to the University for integration testing.

Completion of this milestone starts the three year maintenance clock on the TLCC07 Phase 1 Option clusters.

This milestone is complete when: 1) all Cluster SU and IBA spine switch and cabling hardware is installed, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, the L1 and L2 IBA infrastructure must be functional); 2) the University confirms that the correct Laboratory software stack (see Sections 2.3.2, 2.3.3, and 2.3.4) is installed on the Cluster; 3) the UPR confirms that the Cluster on-site hardware maintenance parts cache is fully stocked; 4) the UPR confirms that the Cluster meets the integration test entry criteria; 5) the acceptance test is successfully executed on the Cluster; and 6) the Cluster successfully completes the integration test exit criteria; 7) the required documentation is approved.

6.3.9 TLCC07 Phase 2 SU Build (TR-1, January 2008)

The Offeror will build, fully assemble, configure, burn-in and test the X SU for TLCC07 Phase 2 defined in section 6.3, as bid, with Tri-Laboratory Linux Build Images as directed by the UPR.

Offeror shall burn in and stress test TLCC07 Phase 2 equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 2 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC07 Phase 2 SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the University confirms that the correct BASS Build Image is installed on TLCC07 Phase 2 SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 2 SU meets the SWL pre-ship test entry criteria; 4) SWL pre-ship test is successfully executed on the TLCC07 Phase 2 SU; 5) The TLCC07 Phase 2 SU successfully completes the SWL pre-ship test exit criteria; 6) the UPR authorizes shipment of SU to tri-Laboratory sites; and 7) all equipment for this milestone leaves Offeror's integration location; 8) the required documentation is approved.

6.3.10 TLCC07 Phase 2 SU Delivery and Acceptance (TR-1, January 2008)

Offeror will deliver the TLCC07 SU to the Tri-Laboratory sites as directed by the UPR, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over the TLCC07 SU to the Tri-Laboratory community for acceptance testing. In addition, Offeror will deliver sufficient IBA spine switches and cables, as directed by the UPR, in order to allow the Tri-Laboratory community to assemble multiple SU clusters.

Offeror will burn in and stress test TLCC07 Phase 2 equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 2 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC07 Phase 2 SU with IBA hardware is installed at the tri-Laboratory sites, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the UPR confirms that the correct BASS Build Image is installed on the TLCC07 Phase 2 SUs; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 2 SUs meet the SWL post-ship test entry criteria; 4) SWL post-ship test is successfully executed on the TLCC07 Phase 2 SUs; 5) the TLCC07 Phase 2 SUs successfully completes the SWL post-ship test exit criteria; 6) the required documentation is approved.

6.3.11 TLCC07 Phase 2 Cluster Integration (TR-1, February 2008)

Offeror will deliver Phase 2 on-site hardware maintenance parts cache to each Phase 2 tri-Laboratory site. Offeror will integrate the site specific SU's into the single multi-SU (2, 4, or 8 SU) fully functional Clusters. The Offeror will fully assemble, pass University's delivery check list and initial functionality and performance verification testing, and turn over the integrated multi-SU Clusters to the University for integration testing.

Completion of this milestone starts the three year maintenance clock on the TLCC07 Phase 2 clusters.

This milestone is complete when: 1) all Cluster SU and IBA spine switch and cabling hardware is installed, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, the L1 and L2 IBA infrastructure must be functional); 2) the University confirms that the correct Laboratory software stack (see Sections 2.3.2, 2.3.3, and 2.3.4) is installed on the Cluster; 3) the UPR confirms that the Cluster on-site hardware maintenance parts cache is fully stocked; 4) the UPR confirms that the Cluster meets the integration test entry criteria; 5) the acceptance test is successfully executed on the Cluster; and 6) the Cluster successfully completes the integration test exit criteria; 7) the required documentation is approved..

6.3.12 TLCC07 Phase 3 Build (TR-1, April 2008)

The Offeror will build, fully assemble, configure, burn-in and test the Y SU for TLCC07 Phase 3 defined in section 6.3, as bid, with Tri-Laboratory Linux Build Images as directed by the UPR.

Offeror shall burn in and stress test TLCC07 Phase 3 equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 3 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC07 Phase 3 SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be

functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the University confirms that the correct BASS Build Image is installed on TLCC07 Phase 3 SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 3 SU meets the SWL pre-ship test entry criteria; 4) SWL pre-ship test is successfully executed on the TLCC07 Phase 3 SU; 5) The TLCC07 Phase 3 SU successfully completes the SWL pre-ship test exit criteria; 6) the UPR authorizes shipment of SU to tri-Laboratory sites; and 7) all equipment for this milestone leaves Offeror's integration location; 8) the required documentation is approved.

6.3.13 TLCC07 SU Phase 3 Delivery and Acceptance (TR-1, April 2008)

Offeror will deliver the TLCC07 SU to the Tri-Laboratory sites as directed by the UPR, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over the TLCC07 SU to the Tri-Laboratory community for acceptance testing. In addition, Offeror will deliver sufficient IBA spine switches and cables, as directed by the UPR, in order to allow the Tri-Laboratory community to assemble multiple SU clusters.

Offeror will burn in and stress test TLCC07 Phase 3 equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 3 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC07 Phase 3 SU with IBA hardware is installed at the tri-Laboratory sites, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the UPR confirms that the correct BASS Build Image is installed on the TLCC07 Phase 3 SUs; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 3 SUs meet the SWL post-ship test entry criteria; 4) SWL post-ship test is successfully executed on the TLCC07 Phase 3 SUs; 5) the TLCC07 Phase 3 SUs successfully completes the SWL post-ship test exit criteria; 6) the required documentation is approved.

6.3.14 TLCC07 Phase 3 Cluster Integration (TR-1, May 2008)

Offeror will deliver Phase 3 on-site hardware maintenance parts cache to each Phase 3 tri-Laboratory site. Offeror will integrate the site specific SU's into the single multi-SU (2, 4, or 8 SU) fully functional Clusters. The Offeror will fully assemble, pass University's delivery check list and initial functionality and performance verification testing, and turn over the integrated multi-SU Clusters to the University for integration testing.

This milestone is complete when: 1) all Cluster SU and IBA spine switch and cabling hardware is installed, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, the L1 and L2 IBA infrastructure must be functional); 2) the

University confirms that the correct Laboratory software stack (see Sections 2.3.2, 2.3.3, and 2.3.4) is installed on the Cluster; 3) the UPR confirms that the Cluster on-site hardware maintenance parts cache is fully stocked; 4) the UPR confirms that the Cluster meets the integration test entry criteria; 5) the acceptance test is successfully executed on the Cluster; and 6) the Cluster successfully completes the integration test exit criteria; 7) the required documentation is approved..

Offeror shall burn in and stress test TLCC07 Phase 3 equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 3 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

Completion of this milestone starts the three year maintenance clock on the TLCC07 Phase 3 clusters.

This milestone is complete when: 1) the TLCC07 Phase 3 SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the University confirms that the correct BASS Build Image is installed on TLCC07 Phase 3 SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 3 SU meets the SWL pre-ship test entry criteria; 4) SWL pre-ship test is successfully executed on the TLCC07 Phase 1 Option SU; 5) The TLCC07 Phase 3 SU successfully completes the SWL pre-ship test exit criteria; 6) the UPR authorizes shipment of SU to tri-Laboratory sites; and 7) all equipment for this milestone leaves Offeror's integration location; 8) the required documentation is approved.

6.3.15 TLCC07 SU Phase 3 Option Delivery and Acceptance (TR-1, October 2007)

Offeror will deliver the TLCC07 SU to the Tri-Laboratory sites as directed by the UPR, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over the TLCC07 SU to the Tri-Laboratory community for acceptance testing. In addition, Offeror will deliver sufficient IBA spine switches and cables, as directed by the UPR, in order to allow the Tri-Laboratory community to assemble multiple SU clusters.

Offeror will burn in and stress test TLCC07 Phase 3 Option equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any

software modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 3 Option cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC07 Phase 3 Option SU with IBA hardware is installed at the tri-Laboratory sites, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the UPR confirms that the correct BASS Build Image is installed on the TLCC07 Phase 3 Option SUs; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 3 Option SUs meet the SWL post-ship test entry criteria; 4) SWL post-ship test is successfully executed on the TLCC07 Phase 3 Option SUs; 5) the TLCC07 Phase 3 Option SUs successfully completes the SWL post-ship test exit criteria; 6) the required documentation is approved.

6.3.16 TLCC07 Phase 3 Option Cluster Integration (TR-1, November 2007)

Offeror will deliver Phase 3 Option on-site hardware maintenance parts cache to each Phase 3 Option tri-Laboratory site. Offeror will integrate the site specific SU's into the single multi-SU (2, 4, or 8 SU) fully functional Clusters. The Offeror will fully assemble, pass University's delivery check list and initial functionality and performance verification testing, and turn over the integrated multi-SU Clusters to the University for integration testing.

Completion of this milestone starts the three year maintenance clock on the TLCC07 Phase 3 option clusters.

This milestone is complete when: 1) all Cluster SU and IBA spine switch and cabling hardware is installed, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, the L1 and L2 IBA infrastructure must be functional); 2) the University confirms that the correct Laboratory software stack (see Sections 2.3.2, 2.3.3, and 2.3.4) is installed on the Cluster; 3) the UPR confirms that the Cluster on-site hardware maintenance parts cache is fully stocked; 4) the UPR confirms that the Cluster meets the integration test entry criteria; 5) the acceptance test is successfully executed on the Cluster; and 6) the Cluster successfully completes the integration test exit criteria; 7) the required documentation is approved.

6.3.17 TLCC07 Phase 4 Build (TR-1, July 2008)

The Offeror will build, fully assemble, configure, burn-in and test the Y SU for TLCC07 Phase 4 defined in section 6.3, as bid, with Tri-Laboratory Linux Build Images as directed by the UPR.

Offeror shall burn in and stress test TLCC07 Phase 4 equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software

modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 4 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC07 Phase 4 SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the University confirms that the correct BASS Build Image is installed on TLCC07 Phase 4 SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 4 SU meets the SWL pre-ship test entry criteria; 4) SWL pre-ship test is successfully executed on the TLCC07 Phase 4 SU; 5) The TLCC07 Phase 4 SU successfully completes the SWL pre-ship test exit criteria; 6) the UPR authorizes shipment of SU to tri-Laboratory sites; and 7) all equipment for this milestone leaves Offeror's integration location; 8) the required documentation is approved.

6.3.18 TLCC07 SU Phase 4 Delivery and Acceptance (TR-1, July 2008)

Offeror will deliver the TLCC07 SU to the Tri-Laboratory sites as directed by the UPR, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over the TLCC07 SU to the Tri-Laboratory community for acceptance testing. In addition, Offeror will deliver sufficient IBA spine switches and cables, as directed by the UPR, in order to allow the Tri-Laboratory community to assemble multiple SU clusters.

Offeror will burn in and stress test TLCC07 Phase 4 equipment (including IBA 4x DDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x DDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48 hour stress test shall be approved by the UPR. Offeror shall demonstrate that the IBA 4x DDR interconnect with this TLCC07 Phase 4 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC07 Phase 4 SU with IBA hardware is installed at the tri-Laboratory sites, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x DDR L1 infrastructure must be functional); 2) the UPR confirms that the correct BASS Build Image is installed on the TLCC07 Phase 4 SUs; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) the University confirms that TLCC07 Phase 4 SUs meet the SWL post-ship test entry criteria; 4) SWL post-ship test is successfully executed on the TLCC07 Phase 4 SUs; 5) the TLCC07 Phase 4 SUs successfully completes the SWL post-ship test exit criteria; 6) the required documentation is approved.

6.3.19 TLCC07 Phase 4 Cluster Integration (TR-1, August 2008)

Offeror will deliver Phase 4 on-site hardware maintenance parts cache to each Phase 4 tri-Laboratory site. Offeror will integrate the site specific SU's into the single multi-SU (2, 4,

or 8 SU) fully functional Clusters. The Offeror will fully assemble, pass University's delivery check list and initial functionality and performance verification testing, and turn over the integrated multi-SU Clusters to the University for integration testing.

Completion of this milestone starts the three year maintenance clock on the TLCC07 Phase 4 clusters.

This milestone is complete when: 1) all Cluster SU and IBA spine switch and cabling hardware is installed, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, the L1 and L2 IBA infrastructure must be functional); 2) the University confirms that the correct Laboratory software stack (see Sections 2.3.2, 2.3.3, and 2.3.4) is installed on the Cluster; 3) the UPR confirms that the Cluster on-site hardware maintenance parts cache is fully stocked; 4) the UPR confirms that the Cluster meets the integration test entry criteria; 5) the acceptance test is successfully executed on the Cluster; and 6) the Cluster successfully completes the integration test exit criteria; 7) the required documentation is approved..

End of Section 6

7 Glossary

7.1 General

Mandatory requirements designated as (MR)	Mandatory requirements (designated MR) are performance features that are essential to Tri-Laboratory requirements, and an Offeror must satisfactorily propose all Mandatory Requirements in order to have its proposal considered responsive.
Target Requirements designated as (TR-1, TR-2 and 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 Tri-Laboratory objectives, but will not result in a nonresponsive determination if omitted from a proposal. Target Requirements add value to a proposal. Target Requirements are prioritized by dash number. TR-1 is most desirable to the University, while TR-2 is more desirable than TR-3. Target Requirements responses will be considered as part of the proposal evaluation process.

7.2 Hardware

b	bit. A single, indivisible binary unit of electronic information.
B	Byte. A collection of eight (8) bits.
32b floating-point arithmetic	Executable binaries (user applications) with 32b (4B) floating-point number representation and arithmetic. Note that this is independent of the number of bytes (4 our 8) utilized for memory reference addressing.
32b virtual memory addressing	All virtual memory addresses in a user application are 32b (4B) integers. Note that this is independent of the type of floating-point number representation and arithmetic.
64b floating-point arithmetic	Executable binaries (user applications) with 64b (8B) floating-point number representation and arithmetic. Note that this is independent of the number of bytes (4 our 8) utilized for memory reference addressing.
64b virtual memory addressing	All virtual memory addresses in a user application are 64b (8B) integers. Note that this is independent of the type of floating-point number representation and arithmetic. Note that all user applications should be compiled, loaded with Offeror supplied libraries and executed with 64b virtual memory addressing by default.
CE	On-site hardware customer engineer performing hardware installation or maintenance (with DOE P-clearance for LLNL).
Cluster	A set of SMPs connected via a scalable network technology. The network will support high bandwidth, low latency message passing. It will also support remote memory referencing.
CPU or core or processor	Central Processing Unit or “core” or processor. A VLSI chip constituting one or more computational core(s) (integer, floating point, and branch units), registers and memory interface (virtual memory translation, TLB, and bus controller) and associated cache.
FLOP or OP	Floating Point OPeration.
FLOPS or OPS	Plural of FLOP.

FLOP/s or OP/s	Floating Point OPeration per second.
FRU	Field Replaceable Unit (FRU) is an aggregation of parts that is a single unit and can be replaced upon failure in the field.
FSB	Front-side bus
GB	gigaByte. gigaByte is a billion base 10 bytes. This is typically used in every context except for Random Access Memory size and is 10^9 (or 1,000,000,000) bytes.
GiB	gibiByte. gibiByte is a billion base 2 bytes. This is typically used in terms of Random Access Memory and is 2^{30} (or 1,073,741,824) bytes. For a complete description of SI units for prefixing binary multiples see URL: http://physics.nist.gov/cuu/Units/binary.html.
GFE	Government Furnished Equipment (GFE) is equipment supplied to the Offeror by the Tri-Laboratory's when TLCC SU build or installation takes place.
GFLOP/s or GOP/s	gigaFLOP/s. Billion (10^9=1,000,000,000) 64-bit floating point operations per second.
HSC	Hot Spare Cluster. A set of nodes on-site at LLNL, LANL and SNL that can be used as a hot spare pool constructed as a stand alone cluster. This HSC will be used to run diagnostics on failing nodes (after they are swapped out of TLCC) to determine root cause for failures and to potentially test software releases.
ISA	Instruction Set Architecture
MB	megaByte. megaByte is a million base 10 bytes. This is typically used in every context except for Random Access Memory size and is 10^6 (or 1,000,000) bytes.
MiB	mebiByte. mebiByte is a million base 2 bytes. This is typically used in terms of Random Access Memory and is 2^{20} (or 1,048,576) bytes. For a complete description of SI units for prefixing binary multiples see URL: http://physics.nist.gov/cuu/Units/binary.html
MDS	Lustre Meta Data Server. Performs the Lustre file system functions associated with file system layout and name space mapping.
MFLOP/s or MOP/s	megaFLOP/s. Million (10^6=1,000,000) 64-bit floating point operations per second.
MTBF	Mean Time Between Failure. A measurement of the expected reliability of the system or component. The MTBF figure can be developed as the result of intensive testing, based on actual product experience, or predicted by analyzing known factors. See URL: http://www.t-cubed.com/faq_mtbf.htm
Node	Four socket AMD x86-64 or Intel EM64T (or binary compatable) quad core die in an SMP configuration with the Linux operating system and IBA HCA.
OSS	Lustre Object Storage Server. The hardware and software associated with the Lustre Object Storage Targets. OSS connect to TLCC via 10 Gb/s Ethernet.
PDU	Power Distribution Unit. Mechanism by which power is distributed to nodes from the higher amperage wall panel.

Peak Rate	The maximum number of 64-bit floating-point instructions (add, subtract, multiply or divide) per second that could conceivably be retired by the system. For microprocessors the peak rate is typically calculated as the maximum number of floating point instructions retired per clock times the clock rate.
POST	Power-On Self Test (POST) is a set of diagnostics that run when the node is powered on to detect all hardware components and verify correct functioning.
SPC	Serial Port Concentrator (SPC) is a rack mounted device (that may be combined with the RPC) that connects the serial ports of nodes to the management Ethernet via reverse telnet protocol. This allows system administrators to log into the serial port of every node via the management network and perform management actions on the node. In addition, this interface allows the system administrators to set up telnet sessions with each node and log all console traffic.
Scalable	A system attribute that increases in performance or size as some function of the peak rating of the system. The scaling regime of interest is at least within the range of 1 teraFLOP/s to 60.0 (and possibly to 120.0) teraFLOP/s peak rate.
SMP	Shared memory Multi-Processor. A set of CPUs sharing random access memory within the same memory address space. The CPUs are connected via a high speed, low latency mechanism to the set of hierarchical memory components. The memory hierarchy consists of at least processor registers, cache and memory. The cache will also be hierarchical. If there are multiple caches, they will be kept coherent automatically by the hardware. The main memory will be UMA architecture. The access mechanism to every memory element will be the same from every processor. More specifically, all memory operations are done with load/store instructions issued by the CPU to move data to/from registers from/to the memory.
SU	Scalable Unit (SU) is the (nearly) identical replicate unit of hardware envisioned by this statement of work.
Tera-Scale	The environment required to fully support production-level, realized teraFLOP/s performance. This environment includes a robust and balanced processor, memory, mass storage, I/O, and communications subsystems; robust code development environment, tools and operating systems; and an integrated cluster wide systems management and full system reliability and availability.
TB	TeraByte. TeraByte is a trillion base 10 bytes. This is typically used in every context except for Random Access Memory size and is 10^{12} (or 1,000,000,000,000) bytes.
TiB	TebiByte. TebiByte is a trillion bytes base 2 bytes. This is typically used in terms of Random Access Memory and is 2^{40} (or 1,099,511,627,776) bytes. For a complete description of SI units for prefixing binary multiples see URL: http://physics.nist.gov/cuu/Units/binary.html
TFLOP/s	teraFLOP/s. Trillion (10^{12}=1,000,000,000,000) 64-bit floating point operations per second.

UMA	Uniform Memory Access architecture. The distance in processor clocks between processor registers and every element of main memory is the same. That is, a load/store operation has the same latency, no matter where the target location is in main memory.
UFGP	University Furnished Government Property.

7.3 Software

32b executable	Executable binaries (user applications) with 32b (4B) virtual memory addressing. Note that this is independent of the number of bytes (4 or 8) utilized for floating-point number representation and arithmetic.
64b executable	Executable binaries (user applications) with 64b (8B) virtual memory addressing. Note that this is independent of the number of bytes (4 or 8) utilized for floating-point number representation and arithmetic. Note that all user applications should be compiled, loaded with Offeror supplied libraries and executed with 64b virtual memory addressing by default.
API (Application Programming Interface)	Syntax and semantics for invoking services from within an executing application. All APIs will be available to both Fortran and C programs, although implementation issues (such as whether the Fortran routines are simply wrappers for calling C routines) are up to the supplier.
BASS	Tri-Laboratory Build and Accept Software Stack (BASS) will be used by the Offeror to build, debug, pre-ship test and accept the TLCC SU hardware.
BIOS	Basic Input-Output System (BIOS) is low level (typically assembly language) code usually held in flash memory on the node that tests and functions the hardware upon power-up or reset or reboot and loads the operating system.
Current standard	Term applied when an API is not "frozen" on a particular version of a standard, but will be upgraded automatically by Offeror as new specifications are released (e.g., "MPI version 2.0" refers to the standard in effect at the time of writing this document, while "current version of MPI" refers to further versions that take effect during the lifetime of this subcontract.
EDAC	Error Detection and Correction (EDAC) software based on the BlueSmoke technology (http://www.sourceforge.net/projects/bluesmoke/)
Fully supported (as applied to system software and tools)	A product-quality implementation, documented and maintained by the HPC machine supplier or an affiliated software supplier.
Gang Scheduling	When a user job is scheduled to run, the Gang scheduler must contemporaneously allocate to CPUs all the threads and processes within that job (either within an SMP or within the cluster of SMPs). This scheduling capability must control all threads and processes within the SMP cluster environment.
GFS	Government Furnished Software (GFS) is software supplied to the Offeror by the University when TLCC build or installation takes place.

Job	A job is a cluster wide abstraction similar to a POSIX session, with certain characteristics and attributes. Commands will be available to manipulate a job as a single entity (including kill, modify, query characteristics, and query state). The characteristics and attributes required for each session type are as follows: 1) interactive session: an interactive session will include all cluster wide processes executed as a child (whether direct or indirect through other processes) of a login shell and will include the login shell process as well. Normally, the login shell process will exist in a process chain as follows: init, inetd, [sshd | telnetd | rlogind | xterm | cron], then shell. 2) batch session: a batch session will include all cluster wide processes executed as a child (whether direct or indirect through other processes) of a shell process executed as a child process of a batch system shepherd process, and will include the batch system shepherd process as well. 3) ftp session: an ftp session will include an ftpd and all its child processes. 4) kernel session: all processes with a pid of 0. 5) idle session: this session does not necessarily actually consist of identifiable processes. It is a pseudo-session used to report the lack of use of resources. 6) system session: all processes owned by root that are not a part of any other session.
LinuxBIOS	An implementation of Linux stored in the node BIOS. This allows nodes to boot from BIOS flash ROM in less than thirty seconds. See http://www.linuxbios.org.
Lustre	Lustre is an open source cluster wide file system based on object technology. See www.lustre.org for more details. SU delivered to LLNL and Sandia will be configured with hardware and software to provide a Lustre cluster wide file system.
MPI	Message Passing Interface Version 1.2 or later. See, for example, http://www-unix.mcs.anl.gov/mpi/mpich/, or http://www.mpi-forum.org/docs/mpi-20-html/mpi2-report.html
OSPF	Open Shortest Path First protocol. See, for example, http://www.ietf.org/rfc/rfc2328.txt
Panasas PanFS	PanFS is a proprietary cluster wide file system hardware and software solution based on industry standard object and interface specifications. See www.panasas.com for more details. SU delivered to LANL will be configured with hardware and software to provide a Panasas cluster wide file system.
Published (as applied to APIs):	Where an API is not required to be consistent across platforms, the capability lists it as "published," referring to the fact that it will be documented and supported, although it will be Offeror- or even platform-specific.
Single-point control (as applied to tool interfaces)	Refers to the ability to control or acquire information on all processes/PEs using a single command or operation.
Standard (as applied to APIs)	Where an API is required to be consistent across platforms, the reference standard is named as part of the capability. The implementation will include all routines defined by that standard (even if some simply result in no-ops on a given platform).

XXX-compatible (as applied to system software and tool definitions)	Requires that a capability be compatible, at the interface level, with the referenced standard, although the lower-level implementation details will differ substantially (e.g., "NFSv4-compatible" means that the distributed file system will be capable of handling standard NFSv4 requests, but need not conform to NFSv4 implementation specifics).

End of Section 7