REQUEST FOR PROPOSAL
FOR

HIGH PERFORMANCE COMPUTING (HPC) SOLUTION

Ref. No. PHY/ALK/43 (27/11/2012)

by
DEPARTMENT OF PHYSICS
UNIVERSITY OF PUNE
PUNE - 411 007
INDIA

NOVEMBER 27, 2012

1 Purpose of this Document

The Department of Physics, University of Pune, plans to purchase a high-performance computing (HPC) cluster system for the needs of in-house scientific research in the area of computational material sciences and soft condense matter. We are looking for a turn-key end-to-end solution complete with hardware, software, implementation and support. The purpose of this Request for Proposal (RFP) is to describe in detail (a) our requirements and (b) our terms and conditions.

2 Terms and Conditions

2.1 Budgetary Provision for this purchase

Rs. 48,00,000/-

2.2 Validity of Pricing

Minimum 120 days from the last date of proposal submission.

2.3 Peak Performance of Cluster

Cluster should have a peak throughput of at least 3.5 TFlops (±5%) with 75% efficiency, excluding the master node (the number of Compute nodes should be configured to match this peak throughput value)

2.4 Vendor Eligibility and Representation

Original Equipment Manufacturer (OEM) of HPC system may take support from System Integrators (SI). The final sales, support and services must be provided directly from OEM.

1. **Vendor Eligibility**: A Vendor must satisfy the following requirements to be eligible to submit proposals.

   (a) International original equipment manufacturers (OEM) with proven track-record in building and supporting HPC cluster platforms for scientific research.

   (b) Continuous presence on `http://www.top500.org/` during 2007-12, with minimum 5 entries in the latest release of November 2012 list.

   (c) Adequate documented experience during 2007-12 in setting up HPC clusters capable of at least 5 TFlops sustained performance.

   (d) Adequate support infrastructure in India (preferably in the Pune region).

   (e) Adequate representation in India's scientific establishments.

2. **Vendor Representation**: An eligible vendor may designate, at their discretion and convenience, only one SI's as a designated partner,

   (a) to represent the vendor during the purchase process,

   (b) for implementing the HPC cluster system and

   (c) for providing all post-implementation service and support.

   Vendor with more than one partner will be disqualified.

3. **Designated partner's / SI's Eligibility**: A partner must satisfy the following requirements to be eligible

   (a) Designated partner must be provided with explicit authorization letters by the vendor.

   (b) Partner should have experience of in implementing and supporting minimum 3 Linux based HPC cluster systems in scientific establishments in India.

   (c) Out of these three HPC systems, at least one must be of size 4TF (peak) or more and remaining two must be of size of at lest 2TF (peak) or above, in government organizations. Partner must have to provide relevant PO (which must be in the name of partner), contract, installation report.

   (d) Partner must have a sales and service office in Pune

4. **Single-Point-of-Contact Support**: We require a single point of contact with the vendor for purchase process, implementation, post-implementation and warranty support, irrespective of who represents the vendor.

2.5 Proposal Evaluation and Vendor Selection

1. All submitted proposals will be screened by a Technical Committee for their technical merit relative to the needs of proposed scientific research, computing power, power and cooling requirements, etc., and will be ranked accordingly. Technical evaluation of submitted proposals will be based on (a) benchmark results (see Sec. 5 for details), and (b) hardware configuration offered (see Sec. 4 for minimum requirements). Submission of benchmark results is a must for this purchase. Hardware configuration/s offered must satisfy the minimum requirements of Sec. 4. Proposals offering a technically superior hardware configuration with greater computing power (as measured by the number of compute nodes) will be given preferential treatment. In all technical matters, the decision of the Technical Committee will be final.
2. Qualified technical proposals will be screened for commercials. Vendors will be called for clarification if required.
3. Final pricing negotiations and vendor selection will take place in a University of Pune Purchase Committee meeting. The date of this meeting will be communicated to qualified vendors.
4. University of Pune reserves the right to disqualify any or all proposals without giving any reasons. University of Pune is not bound to necessarily accept the lowest-priced proposal.

2.6 Delivery

A purchase order will be issued by the University to the vendor selected by the Purchase Committee. We expect delivery of the HPC cluster system in its entirety within 4/6 weeks after the date of this purchase order.

2.7 Implementation

We expect implementation of the HPC cluster system to be completed by the vendor within 2 weeks after delivery. End-goals of implementation are: (a) the deployment of the HPC cluster system complete with hardware, OS and clusterware, and user-specified software, and (b) a clear demonstration that the system is fully functional and usable for the end-user for scientific/computational research.

2.8 Testing and Certification

The warranty on the HPC cluster system will begin on the date the HPC cluster is demonstrated by the vendor to the Department of Physics's technical team to be fully operational and working satisfactorily. This date will be decided as follows: Upon completion of implementation of the entire cluster system (hardware+software) by the vendor, the departmental technical team will test it for not more than one week at full computational load. If no problems of any kind show up during this test period, the system will be certified by the departmental technical team as "fully functional and working satisfactorily". Warranty on the cluster system will begin on the day of this certification.

If any problems show up, they will need to be corrected by the vendor and the departmental technical team will again subject the cluster system through the mandatory testing period. This test cycle will be repeated as many times as required until the cluster system is demonstrated to be fully functional to the departmental technical team's satisfaction.

2.9 Warranty and Support

1. **Warranty**. Your proposal must provide, in the least,
   (a) 3-year on-site comprehensive warranty with next-business-day response/support for all hardware.
   (b) 3-year on-site next-business-day support for everything related to the operating system, clusterware and software setup.
2. **Single-Point-of-Contact Support**. Irrespective of who represents the vendor (see Sec. 2.4), we need one single point of contact with the vendor for all and post-implementation/warranty support.

2.10 Payment

Upon certification of the fully-implemented HPC cluster system by the departments technical team as "fully- operational and working satisfactorily (see Sec. 2.8 for details), University of Pune will make full payment within 4 weeks by a mutually agreeable method.

3 Proposal Submission

3.1 General Instructions

1. **Technical and Commercial Proposals**. Technical and commercial proposals should be addressed to The Head, Department of Physics, University of Pune, Pune-411 007, INDIA and submitted in separate sealed envelopes. Envelopes containing proposals should clearly indicate
   (a) vendor name,
   (b) type of proposal: technical or commercial and
   (c) our reference number for this Tender Notice/RFP (see below).

   Detailed instructions on the format of the technical and commercial proposals can be found in Sec. 3.2.

2. **Our Reference Number**. Please quote our reference number "PHY/ALK/43 (27/11/2012)" on all correspondence.
3. **Clarification Enquiries**. Any clarification enquiries may be directed to the Prof. Anjali Kshirsagar, Department of Physics, University of Pune, Pune 411 007 India Phone: (20).2560.1412. Extn. 317 Email: anjali@physics.unipune.ac.in.
4. **Last Date for Proposal Submission: December 4, 2012 at 04:00 PM.**

3.2 Response Format

3.2.1 Technical Proposal

1. **Vendor Information.** Vendor profile, together with sufficiently detailed notes on
   (a) Expertise+experience in building HPC cluster systems with Linux- based software setup.
   (b) Support infrastructure in India and in the Pune region.
   (c) Presence in Indias scientific establishments: Provide sufficiently detailed information including name of the establishment, purpose of the HPC cluster system supplied, nature/configuration of the HPC cluster system, year of purchase, contact person information if available.
   (d) Any prior presence on the University of Pune campus.
2. **Vendor Representation.** Sufficiently detailed information and profile of the implementation and support partners, if applicable (see Sec. 2.4), together with contact details and focus on their experience in implementing and supporting Linux-based HPC cluster systems in scientific establishments in India.
3. **Complete Technical Specification of the Offered HPC Cluster System.** Technical specifications of the offered system must satisfy the minimum requirements of Sec. 4. **A vendor may offer more than one alternative/option for the hardware configuration.** For each hardware configuration alternative offered, we need the following information:
   (a) Complete spec-sheets, brochures, and URLs to information pages on the vendor's website, etc., for each major cluster component (compute nodes, master node, GPU, any peripherals, software setup, etc.).
   (b) Peak power rating of the complete system, plus power and cooling requirements.
   (c) A clear summary of what is offered over and above the minimum requirements.
4. **Results of the Benchmark Tests.** Include a CD with requested results of benchmark tests (Sec. 5) together with relevant information about the hardware configuration + software set-up of the system used for running the benchmarks.

3.2.2 Commercial Proposal

1. **Vendor Representation.** Relevant information and profile of the business partner, if applicable (see Sec. 2.4), together with contact details.
2. Complete Pricing Details of the Offered HPC Cluster System. For each of the technical alternatives proposed:

(a) Pricing for each major HPC cluster system component (master node, compute nodes, infiniband switch, GPU, software and implementation).

(b) Per-node pricing of the offered compute node.

All hardware prices are to be quoted in USD (CIF/CIP Mumbai): Customs clearance will be taken care of by the University of Pune. University of Pune is exempted from octroi duties levied by the Pune Municipal Corporation and the department's office will provide an octroi exemption certificate if and when necessary. Any other applicable taxes should be mentioned clearly.

3. **Warranty Details.** We assume that all terms and conditions from our side (see Sec. 2) are accepted by the vendor. Any additional features offered over and above the minimum required warranty and support terms (Sec. 2.9) should be clearly mentioned.
4. **Information about Single-Point-of-Contact for Warranty and Support.** Complete contact information for the single-point-of-contact for warranty support (Sec. 2.9).

4 Minimum Requirements for the HPC Cluster System

We require a HPC cluster system consisting of

1. One master node,
2. At least 9 standard rack mount servers (Not more than 2U size) to act as compute nodes, Out of these 9 compute nodes,
   (a) one node should have 8GB/core RAM
   (b) one node should have two latest NVIDIA GPU cards (Preferably Tesla M2090) in 1:1 (CPU:GPU) ratio
3. Infiniband QDR 18 ports Switch as primary interconnect the nodes with appropriate connecting cables
4. A 24 port gigabit Ethernet switch as Secondary interconnect with connecting cables for Cluster Administration, Rack Mount Unmanaged Switch
5. Keyboard, Mouse, 19 inch LCD monitor and 16 port KVM (D-LINK) switch for central monitoring and management
6. An internal storage (to be physically added in the master node) with at least 5 TB of usable storage at RAID 5,
7. Software components (operating system, cluster management and monitoring software, compilers, etc.) necessary for the operation of the complete system as a scientific computing platform, plus implementation and deployment of the complete system.
8. Appropriate quantity of PDU's (16 A or higher) and Cable managers as required. All PDU's should be compatible with quoted servers.

Minimum requirements on each of these components are specified in Sec. 4.1-4.6.

4.1 The Master Node

Configuration	The master node should have the same base-level configuration as for a compute node, except for the form factor.
Storage	The master should have 5 TB of usable storage using SATA HDD.
Form Factor	Form factor in case of master should be 2U.
Quantity	1

4.2 Compute Nodes

Processor	Two Intel Sandy Bridge eight core CPU preferably 2.6GHz
L3 Cache	20MB.
Chipset	Intel C600 chipset
Memory	64 GB ECC DDR3 Memory 1666 MHz and scalable up to 256GB memory
DIMM Slots	Should support 20 DIMM Slots memory configurations
Memory Property	Should support ECC,Chipkill, memory mirroring and memory sparing.
Hard disk Drives	Server should be configured with 500GB SATA HDD. Should support intermix of SAS and SATA.
RAID Controller	Hardware Raid controller supporting basic hardware RAID 0, 1 must
Graphics Controller	16MB Memory
Infiniband Ports	Should be configured with IB switch ports
Gigabit Ethernet Ports	Should be configured with Gigabit Ethernet ports
High speed Interconnect	4x QDR Infiniband Dual port Adaptor for building highspeed primary network.
PCI Slots	Support minimum 2 PCIe slots with dual processor configuration.
Ports	Should support 2 Number of external USB ports, also should support Video ports.
Power Supply	Should be configured with power supplies with 80 PLUS Platinum certification .
Management Functionality	Should support integrated management , Server should be supplied with OEM Server Management software .
Form Factor	1U Rack mountable.
Server make	International OEM Make only with at least 5 entry in top500.org as per latest November 2012 list.
OS Support	Latest version of Scientific Linux (SL), Redhat, Suse Linux.
Warranty	3 Years Comprehensive Onsite Warranty .

4.3 Infiniband Switch

Speed Rating	QDR Switch as primary interconnect.
Number of Ports	18(minimum).
Form Factor	Standard 1U. Preferably Voltaire/Qlogic (make)
Quantity	1

4.4 Software and Implementation

General Requirements	The proposed HPC cluster system should be deployed with (a) an open-source Linux-based operating system with adequate device driver support; tools for cluster installation and management that support node-group and repository manager for deploying updates, patches, etc., or for quickly re-imaging new nodes with no interruption in uptime; tools for monitoring cluster health,resource usage; and a job scheduler; (b) compilers, MPI, and code development tools; (c) installation/integration of user-specific scientific applications (see below); (d) integration of all software components so as to make the complete HPC cluster system fully functional and usable (e.g.,integration of the scheduler with MPI, any license managers, etc.).
Operating System	Open source HPC Operating system.
Workload and Cluster Management	Latest stable release of a reputed workload and cluster management software suite. Scheduling and cluster management software should support policy-based workload management, graphical cluster administration interface, monitoring and reporting tools, etc. Open-source software preferred.
Compilers and MPI	IntelR Cluster Studio XE 2013 or better for Linux OS - 1 Users floating Academic with 3 years Support (ESD) - Quantity 1
Scientific Applications	Installation, integration and any performance tuning of (a) standard numerical libraries (BLAS,LAPACK, ATLAS, FFTW), (b) SIESTA, Quantum Espresso and VASP (source codes will be provided by us) and other standard open-source scientific tools/applications in consultation with the Department's technical team.
HPC Cluster Implementation	End-goals of implementation are (a) the deployment of the HPC cluster system complete with hardware, OS, clusterware and user-specified software, so that it is functional and usable for the end-user for scientific/computational research and (b) a clear demonstration of the same.

5 Benchmark Tests for Technical Evaluation

Technical evaluation of a proposed HPC cluster solution will be based on the results of a number of benchmark tests described below in Sec. 5.1-5.3. Our benchmark tests have been designed in consideration of the computational resource usage profile (i.e, heavy number-crunching by and large, with moderate to heavy I/O in spurts) of the actual tools that will be used for the scientific research using the proposed HPC cluster.

The HPC cluster system used for benchmarking must satisfy all other minimum requirements (Sec. 4). The benchmarking must be done with CPU only for all three tools mentioned below. For Sec 5.1 and 5.2 benchmarking using GPU will be preferred.

Specifically, we need the following data for each of the benchmark tests described in Sec. 5.1-5.3:

1. The following results on 16, 32, 64, 128 cores (CPU):
   - Run time, as provided by the Unix time command or equivalent.
   - Output generated by each of the runs (`stdout, stderr`, plus any output files). Please make sure that output files do not get overwritten across runs over different number of cores.
2. Complete and detailed configuration information (hardware configuration + software setup) of the HPC cluster that was used to run these benchmark tests.
3. Output (`stdout+stderr`) of the compilation sequence for the tools below.

5.1 VASP

VASP-5.3 (http://www.vasp.at/), The Vienna Ab initio Simulation Package (VASP) is a computer program for atomic scale materials modeling, e.g. electronic structure calculations and quantum-mechanical molecular dynamics, from first principles.

5.2 Quantum Espresso

Quntum Espresso(http://www.quantum-espresso.org/) is an integrated suite of Open-Source computer codes for electronic-structure calculations and materials modeling at the nanoscale. It is based on density-functional theory, plane waves, and pseudopotentials.

5.3 SIESTA

SIESTA(http://icmab.cat/leem/siesta/) is both a method and its computer program implementation, to perform efficient electronic structure calculations and *ab initio* molecular dynamics simulations of molecules and solids.

P.S.: All input files for above packages are available on the URL:
`http://physics.unipune.ac.in/hpc/`