INDIAN INSTITUTE OF SCIENCE EDUCATION AND RESEARCH PUNE

CLARIFICATION ON TENDER NUMBER - IISER-PUR-1718-16

ITEM DESCRIPTION- PROCUREMENT OF HPC CLUSTER

Refer our Press Tender Notice No.IISER/S&P/20/16 dated 25.2.2017 for procurement of HPC Cluster .Tender Reference Number – IISER-PUR-1718-16.

Pre-Bid meeting was held on March 07th, 2017 at 2.30 PM and minutes of meeting is as under.

At the outset, the Chairman welcomed all the Members and the representative of the Prospective Bidders and briefed in general the scope of the Project and thereafter requested Assistant Registrar (S&P) to brief the vendors on the salient features of the commercial terms and the indenting Officer to read out the clarification sought by the Prospective Bidders and replied thereto as detailed in Annexure -II

The representatives present were satisfied with the replies given and it was informed that the corrections / additons / clarifications given, as discussed during the Pre-Bid Conference would be hosted on the website of IISER Pune and all the Prospective Bidders are required to take cognizance of the proceedings of the Pre-Bid Conference before submitting their bids as stipulated in the Bidding Documents.

Chapter 4 , Schedule of Requirement, Specifications & Allied Technical details has been revised and added as Annexure – IV of this document. The last date for submission of tender has been extended to 18/4/2017 and technical bids will be opened on 19/4/2017 – Time 3.30 PM.The

The Tender Document can be downloaded from Central Public Procurement (CPP) Portal https://eprocure.gov.in/eprocure/app or Institute website www.iiserpune.ac.in and bid is to be submitted online only through the E-procurement portal up to the last date and time of submission of tender. The critical dates are given below:

Sr.No	Particulars	Date	Time
1	Bid Submission Start Date	14- March-2017	14:30 Hrs.
2	Bid Submission Close Date	18-April-2017	13:00 Hrs.
3	Closing date & time for Submission of original EMD & Tender Fee	18 April-2017	13:00 Hrs.
4	Opening of Technical Bids	19-April-2017	15:30 Hrs.

The other terms & conditions of the notice issued on our IISER website www.iiserpune.ac .in will remain unchanged.No more correspondence in this regard will be entertained

The meeting ended with vote of thanks to the Chair

Sd/-

7.3.2017 Assistant Registrar (S&P)

ANNEXURE -II

IISER PUNE

PRE-BID CONFERENCE FOR PROCUREMENT OF HPC CLUSTER

TECHNICAL QUERIES AND CLARIFICATION

TENDER NUMBER - IISER-PUR-1718-16 DATE :7.3.17

S.No	Query/Clarification Sought	Clarification / Amendment
01	For Master node, Management IPMI 2.0 with LAN and KVM-over-LAN support, Please mention if only support is required or license also.	Yes, license is required
02	For Master node, Chassis 2U rack mountable with mounting rails, Please allow 1U - we can provide all other specs	2U/1U both are acceptable.
03	For Master Node, RAID Card, Appropriate RAID Controller with at least 2GB cache, supporting RAID 0, 1, 5 & 6. Is RAID 5 & 6 required for 2 disks scenario? Is 2GB cache required for 2 x 1TB SATA disks?	RAID card is required; battery backup cache is not required.

	If yes, do you require battery backup cache?	
04	For Master Node, Peripherals, Rack mountable TFT with keyboard, mouse and KVM, How many ports on KVM switch are required? Also please specify if the KVM switch required is analogue or digital	Minimum 40 ports are required on the KVM switch. Either analogue or digital.
05	For compute node, management, IPMI 2.0 with LAN and KVM-over-LAN support, Please mention if only support is required or licence also	Yes, license is required.
06	PFS based storage solution, Usable capacity for OSSMin. 65 TB usable RAID6 (Can be either Hardware orSoftware RAID). Preferred disk layout - 9D + 2P. NLSAS disks (12Gbps) can be used. Please allow 8D + 2P, which according to our design gives better performance.	65 TB usable to be provided, 8D+2P is allowed.
07	PFS based storage solution, Usable capacity for MDS; Vendor should use min. two 600GB 15K RPM disks (orbetter). Disk count to be decided by the vendorbased on the requirement for achieving the statedperformance, We understand these are for MDT.	Yes, it is MDT.
08	Miscellaneous, Benchmarking andTechnical Qualificationcriteria, IISER Pune reserves the call for doing benchmarking. Please see the technical qualification for the detailsof benchmark. Additional details regardingbenchmark can be found below. When will the information on Benchmark be	Appended as Chapter 4.

	released?	
09	Miscellaneous, Other conditions, The bidder should ensure that all hardware(H/W) components offered in the Bill ofMaterial should be covered under OEM/bidder support enabling program so as to get back end support / benefits from Principles / OEM/ bidder in terms of Free Software Update Support / Maintenance releases, if any, to a particular Software Version. Please change to - "all hardware (H/W) components offered in the Bill of Material should be covered under OEM/~~bidder~~support enabling program so as to get backend support / benefits from Principles / OEM/~~bidder~~ in terms of Free Software UpdateSupport / Maintenance releases, if any, to aparticular Software Version.	No change in the tender conditions.
10	Miscellaneous,**Acceptance criteria,** After installation is complete, the bidder will carry out acceptance tests. This consists of a 72-hour burn test and running each of the benchmarks (using the same software suite they have used while running the benchmarks for the technical bid) at IISER Pune to demonstrate the performance as claimed in the technical bid. Kindly elaborate the clause "(using the same software suite they have used while running thebenchmarks for the technical bid)" As per our understanding, benchmark is to be run as partof Acceptance Test after the supply. If the benchmarks are to be submitted with bid We can run benchmark on nodes with E5-	The benchmarks have to be run on the same configuration machine.

	2690V4 basednodes. Please allow this.	
11	PFS base storage Solution : Performance, Min. 65 TB usable RAID6 (Can be either Hardware or Software RAID). Preferred disk layout - 9D + 2P. NL SAS disks (12Gbps) can be used. We request IISER to clarify that Min. 4 GB/s should be particularly for Read or write or Combination of both	Write performance of 4 GB/s.
12	Benchmarking and Technical qualification criteria: IISER Pune reserves the call for doing benchmarking. Please see the technical qualification for the details of benchmark. Additional details regarding benchmark can be found below. We request IISER to mention On how many number of nodes OEM/bidder needs to run the benchmark? Please provide more details on what will be the expectation of IISER from Bidder for benchmarking	Appended as Chapter 4.
13	Compute nodes: HDD -1 x 200GB enterprise class SATA SSD with minimum 1.0 dwpd. Request IISER to please clarify that is, it 10disk write per day or 1 disk write per day?	It is 1 dwpd.
14	PFS based storage solution, Intel Manager for Lustre, Separate 1U Rack mount Server should be offered with min. 1 x Intel Xeon E5-2603v4 (6c @ 1.7GHz) / 16GB / 2 x 1 TB SATA disks / 4 x 1G / RPS ,Request IISER to increase the capacity of master node(CPU,RAM,Disk) and install Intel Manager for Luster.	No change in tender specifications.
15	Eligibility Criteria for Bidder: The bidder should have documentary evidence of purchase order for at least one successful installation of distributed memory node cluster, with only CPUs, of at least 30 TF peak performances, in	No change in eligibility criteria.

	premiere institutes of India (IISER/IIT/NIT etc.) in a single order within the last 3 years. A certificate of satisfactory performance from such a customer with the name (specific contact person/researcher), email id and phone number of the customer must be provided, We request you to amend the same clause as per below Eligibility Criteria for Bidder :The bidder should have documentary evidence of purchase order for at least one successful installation of distributed memory node cluster, with only CPUs, of at least 20 TF peak performances, in premiere institutes of India (IISER/IIT/NIT etc.) in a single order within the last 3 years. A certificate of satisfactory performance from such a customer with the name (specific contact person/researcher), email id and phone number of the customer must be provided.	
16	Compute Nodes Processor: 2 x Intel Xeon E5-2680 V4 (Fourteen-core, 2.4GHz, 40M). Total28Cores. Queries: a) Since the cluster is to be used for HPC applications in Molecular Dynamics, Quantum Chemistry domain, the preferred processor for such applications offering the best application price/performance is the Intel Xeon E5-2697A V4 (Sixteen-core, 2.6Ghz, 40M). Total 32-cores. Hence we request IISER to look at this best-in class processor for HPC applications on both Master and compute nodes.	No change in tender specifications.
17	Chassis & Power supply Dense form factor nodes occupying on an	Yes.

	average of 0.5U per node to be proposed. If there are any blank slots in the chassis, this will not be considered as part of the 0.5U calculation. Similarly, if there are any External Power shelves in the solution powering the compute nodes, the space used by these Power Shelves shall not be used for the 0.5U calculation. Queries: a) Since IISER is asking for dense form factor, RSC has a world record in compute density with its dense direct-liquid cooled nodes. We can offer upto 6 nodes in 2U space. Kindly confirm if IISER is open for even denser form factor nodes occupying on an average < 0.5U per node	
18	Kindly confirm the level of power supply redundancy required for compute nodes e.g. N+1, N+2 etc.	N+1
19	Otherconditions: The bidder should also give power, UPS and cooling requirements for the cluster solution along with the proposal. The power and cooling requirements should be calculated at peak usage assuming that all cores in every node are being used. Please provide separate power requirement for cluster and storage. Queries: This is a very subjective statement when comparing various solutions. We request IISER to consider the Power, Space, UPS and cooling requirements for L1 calculation. Each bidder has a different solution and based on their uniqueness (air-flow cooled solution/partial	No change in tender specifications.

	liquid cooling/complete liquid cooling solution) the efficiency and cost varies. Therefore to get a energy-efficient solution, it is recommended that IISER considers the total TCO costs (total Power, space, cooling required etc.) by providing appropriate weightage to each solution in the final L1 calculation	
20	Benchmarking and Technical qualification criteria: We request you to kindly confirm If any benchmarking needs to be submitted along with bid, As we did not find any details of benchmarking in tender document.	Appended as Chapter 4.
21	Acceptance criteria : After installation is complete, the bidder will carry out acceptance tests. This consists of a 72-hour burn test and running each of the benchmarks (using the same software suite they have used while running the benchmarks for the technical bid) at IISER Pune to demonstrate the performance as claimed in the technical bid Request you to elaborate this point.	Appended as Chapter 4.
22	**Performance Benchmarks:** The technical evaluation committee needs to be provided with an evaluation system to carry out performance benchmarks. Please help to clarify this point.	Appended as Chapter 4.
23	Refer you bidder eligibility clause for Installation reference of 30TF Rpeak. We would like to inform you that, we have the installation references as per below. 30TF Rpeak - Signoff in Year 2013 28.8 TF Rpeak – Signoff in Year 2016 Request you to please consider the same &	No change in the eligibility conditions of the tender.

	allowed us for bidding against this RFP.	
24	Number of nodes the benchmark has to be run on?	Both the HPL and VASP benchmarks have to be run on 8 nodes.

ANNEXURE -III

IISER PUNE

PRE-BID CONFERENCE FOR PROCUREMENT OF HPC CLUSTER

COMMERCIAL QUERIES AND CLARIFICATION

TENDER NUMBER - IISER-PUR-1718-16 DATE : 7.3.17

S.No	Query/Clarification Sought	Clarification / Amendment
1	**Delivery Period / Timeliness:** The deliveries & installation must be completed **within 01 Month** after placement of purchase order. The time is the essence of the contract. It is mandatory for the BIDDERs who respond to this bid to meet these expectations, as are tightly linked to IISER, PUNE's plans of completing the project within the time frame. Kindly enhance the delivery & installation time from one month to 8 weeks. There are some essential components of asking solution which take some time for delivery	Delivery and installation has been extended and shall be 6 weeks within opening of LC.

2	With reference to pre bid meeting we had, request you to extend the due date of submission by at least a week, as there are benchmark s involved and which would require more time to plan for same.	The due date of the submission has been extended to 18th April, 2017.

CHAPTER 4

Detailed Specifications for the HPC cluster

Sr. No	Parameters	Values or Notes
Master Node: Quantity 1		
1	**Processor**	2 x Intel Xeon E5-2620 V4 (2.1GHz / 8-Core / 20MB) or better
2	**Memory**	Minimum 64GB (4X16GB) DDR4 Reg. ECC 2400MHz with necessary free slots and provision to add at least another 64GB in future
3	**Internal Storage**	2 x 1000 GB, 7200 rpm, Enterprise SATA, Hot Plug HDD's (RAID 1)
4	**NIC**	Dual Gigabit Ethernet RJ45 ports
5	**Management**	IPMI 2.0 with LAN and KVM-over-LAN support
6	**Exp. Slots**	1x PCI-E 3.0 x16, 2x PCI-E 3.0 x8 (Low-profile slots)
7	**Chassis**	2U rack mountable with mounting rails
8	**RAID Card**	Appropriate RAID Controller with at least 2GB cache supporting RAID 0, 1, 5 & 6
9	**Interconnect**	1x Single-Port EDR/OPA with cables of appropriate length
10	**DVD Writer**	DVDRW drive
11	**Power Supply &Cables**	Power cables should be compatible with the quoted rack. • Redundant and hot pluggable platinum power supply.
12	**Warranty**	3 years on-site comprehensive warranty
13	**Peripherals**	Rack mountable TFT with keyboard, mouse and KVM
Compute Node(s): Quantity 36		
1	**Processor**	2 x Intel Xeon E5-2680 V4 (Fourteen-core, 2.4GHz, 40M). Total 28 Cores.
2	**Memory**	8 x 16GB (Total 128GB) DDR4-2400 ECC REG.(Max 512GB, 8 DIMMs)
3	**HDDs**	1 x 200 GB enterprise class SATA SSD with minimum 1.0 dwpd
4	**Management**	IPMI 2.0 with LAN and KVM-over-LAN support

5	**Interconnect**	1x Single-Port EDR/OPA with cables of appropriate length
6	**Exp. Slots**	1x PCI-E 3.0 x16 (LP) slot
7	**Chassis, Power supply**	Dense form factor nodes occupying on an average of 0.5U per node to be proposed. If there are any blank slots in the chassis, this will not be considered as part of the 0.5U calculation. Similarly, if there are any External Power shelves in the solution powering the compute nodes, the space used by these Power Shelves shall not be used for the 0.5U calculation.
8	**NIC**	Dual Gigabit Ethernet
9	**Warranty**	3 years on-site comprehensive warranty
<u>PFS based storage solution</u>		
1	**PFS**	Intel Enterprise Edition Lustre (IEEL) - 3 years only should be offered with necessary support.
2	**Performance**	Min. 4GB/s write/read using IOR or IOZone
3	**Usable capacity for OSS**	Min. 65 TB usable RAID6 (Can be either Hardware or Software RAID). Preferred disk layout - 9D + 2P. NL SAS disks (12Gbps) can be used.
4	**Usable capacity for MDS**	Vendor should use min. two 600GB 15K RPM disks (or better). Disk count to be decided by the vendor based on the requirement for achieving the stated performance
5	**Node requirement for MDS/OSS**	The appliance should be configured with 2 nodes. Node 1 to be configured for MDS and Node 2 for OSS. In case of any node failure, the services should start from the other surviving node, without any manual intervention.
6	**MDS/OSS Node configuration**	Each of the nodes should be configured with 2 x Intel Xeon E5-2690v4 (or better in terms of cores and frequency)/ 256GB/ necessary Disk controllers for the solution/ 100Gbps EDR or OPA (single port).
7	**Operating System**	CentOS. Vendors are free to offer RHEL/SLES incase CentOS is not supported by the appliance.
8	**Spares**	One hard disk to be kept as cold spare onsite for both Metadata and OSS.
9	**Solution form factor**	Should not exceed 12 U, including the IML Server
10	**Redundancy**	Solution should be configured with No Single Point of Failure including Power Supplies and fans
11	**Intel Manager for Lustre**	parate 1U Rack mount Server should be offered with min. 1 x Intel Xeon E5-2603v4 (6c @ 1.7GHz) / 16GB / 2 x 1 TB SATA disks / 4 x 1G / RPS
<u>Primary Interconnect: EDR Infiniband Switch or OPA</u> **<u>Secondary Interconnect: Gigabit Ethernet Switch</u>**		
1	**Infiniband Switch/OPA**	• EDR 100 Gbps per port or OPA. Backplane bandwidth must correspond to 100% non-blocking architecture.

		• The Infiniband/OPA switch should be manageable, rack mounted, with redundant power supply.
2	**Gigabit Ethernet Switch**	At least 48 ports managed Gigabit Ethernet Switch with required number of cables to connect all nodes. All nodes need to be connected by Gigabit network for remote management hardware as well as for system network. Both the networks should be completely different.
Miscellaneous		
1	**Rack**	42 U rack with required number of PDU's and accessories
2	**OS & Cluster Management Software**	Open Source Linux Cent OS & Cluster Management software. The software should be able to do the following: • Add/Remove nodes from operation. • Allow online replacement of failed nodes. • Perform OS installation of all nodes from head node. • Monitor node hardware stats. • Monitor node performance stats. • All of the above to be performed remotely, using both GUI and CLI interfaces. • Jobs and queue handling software should be installed and should be demonstrated to the client.
3	**Benchmarking and Technical qualification criteria**	IISER Pune reserves the call for doing benchmarking. Please see the technical qualification for the details of benchmark. Additional details regarding benchmark can be found below.
4	**Other conditions**	• The server and all its components should be verified and recommended by the motherboard manufacturer by means of compatibility list. • No on-site assembling / integration is allowed after delivery of the hardware, the nodes/equipment must be tested and factory integrated. Only rack-mounting, installation of OS and applications can be done at site. • The bidder should ensure that all hardware (H/W) components offered in the Bill of Material should be covered under OEM/bidder support enabling program so as to get back end support / benefits from Principles / OEM/ bidder in terms of Free Software Update

		Support / Maintenance releases, if any, to a particular Software Version. • During the warranty period, the complaints should be addressed within 24 hours of receiving the same. In case defective parts needed to be replaced during the warranty period (3 years) it should be done within a period of two working days. If the bidder fails to do so in either of the above two cases, a penalty @ Rs. 5000/- per day delay shall be invoked. The corresponding amount shall be deducted from the PBG. • Full technical details of the proposed compute hardware, interconnects and storage must be provided. All model numbers must be provided. • Full details of the rack, such as dimensions, cabling, PDUs, including model numbers must be provided. • Full technical details of the proposed software including OS, clustering suite, cluster monitoring software must be provided. • The bidder should also give power, UPS and cooling requirements for the cluster solution along with the proposal. The power and cooling requirements should be calculated at peak usage assuming that all cores in every node are being used. Please provide separate power requirement for cluster and storage. • Bidders to include all necessary hardware/software and accessories of the solution to render the solution operational under optimum guaranteed parameters even if any item is not included in the general or detailed specifications but found required for the solution, and the same should be to the satisfaction of the customer. • The item wise rates of each item are to be quoted. IISER Pune reserves the right to vary the quantity at the time of Purchase Order placement. • The bidder should provide schematic diagram of the entire solution. • Any other item needed to complete the offered solution should be provided and indicated separately.
5	**Installation, Warranty and**	**1.** The vendor shall install and configure all

	Training	required hardware and open source software suites, including but not limited to racking and stacking, cluster networking, configuring all nodes, installation of open source compilers and applications, configuration of environment variables and license utility configuration. **2.** The vendor to ensure that the hardware and software components are compatible with each other, and provide necessary cables/wires and any other accessories necessary for connecting the supplied components. Bidder must install the complete system interconnecting all the components above. **3.** One day admin and user training on the cluster usage and administration must be provided. **4.** All standard libraries like BLAS, LAPACK, SCALAPACK, ATLAS, fftw, etc. should be installed along with some application software like VASP,QUANTUM EXPRESSO, GROMACS, and Gaussian etc with open source as well as Intel cluster studio. **5.** All standard Open Source compilers like gfortran, gcc, openmpi etc. Should be installed
6	**Optional: Intel Cluster Studio Single user academic license for a period of 1 year and 3 years options**	To be supplied and installed onto the cluster
7	**Compliance statement**	**a)** Compliance statement needs to be provided by vendors clearly specifying COMPLY/NON COMPLY with remarks of all of the points mentioned above.
8	**Acceptance criteria**	After installation is complete, the bidder will carry out acceptance tests. This consists of a 72-hour burn test and running each of the benchmarks (using the same software suite they have used while running the benchmarks for the technical bid) at IISER Pune to demonstrate the performance as claimed in the technical bid.

Details regarding the benchmarks:

- Please unzip and untar the file and follow the INSTRUCTION in the README file carefully.
- VASP version should be 5.3.3. VASP should be installed and used in double precision using opensource compilers (gcc/openmpi). The VASP Code should be obtained from IISER Pune after

signing the declaration form, and Vendors has to clarify that they do not possess the benchmark licence. VASP INPUT files should not be altered.

- Benchmarking of HPL and VASP should be purely on CPU cores and should be done on the same configurations and specs as asked in this tender and with a minimum of 8 nodes.
- Results should be submitted along with technical bid.
- Log files of benchmark runs in electronic format must be provided. No hard copies will be accepted.
- All the input files will be provided.
- A copy of the runtime environment used for the benchmark should be provided.
- Details of VASP installation used to run the program should be provided.
- The script used to run the job should be provided.
- An undertaking that guarantees the reproduction of the benchmark data on the proposed cluster in our premises after installation must be provided. In case of failing to do so (with maximum 7% deviation from the results submitted in the bid), IISER Pune reserves the right to deduct 10% from the payment.
- Complete hardware and software environment details of the machine on which the benchmarks were carried out, must be provided.
- All cores present in each node must be used in running the benchmarked jobs.
- If the benchmarking was carried out on a system containing Y cores per node, the proposed solution too must contain the same number of cores per node.
- **The benchmarks should be run with turbo mode off.** Vendors can employ architecture-specific code enhancements (for instances compiler flags) during benchmarking. No fast math options are allowed. However, these must be transparently mentioned in their proposals and relevant electronic files (makefiles, etc.) used in the benchmarks must be shared.
- The bidder/OEM can run the benchmark. The bidder/OEM should provide an undertaking certifying that they have done the benchmark themselves.

Technical Qualifying Criteria:

- The bidder should carry out below listed bench mark programs:
  1. HPL benchmark (see http://www.netlib.org/benchmark/hpl/)
  2. VASP (http://www.vasp.at/) benchmark. The benchmark files in tar, zipped form may be obtained from the following weblink: http://www.iiserpune.ac.in/~mukul/VASP_Benchmark.tar.gz

- For the HPL benchmark, we will look into the ratio of the sustained vs. theoretical peak.

- A two step criterion will be used for technically qualifying the vendors.

  (a) A minimum of 73% is acceptable for the HPL benchmarks.
  (b) If (a) is satisfied, the benchmarks for VASP will be used for further considerations. An expert technical committee shall decide the technical eligibility of the bidders on the basis of the correctness of the numerical results and the performance of the benchmarks. The results shall be evaluated

on the basis of scalability as well as in terms of efficiency. However if the submitted benchmark is more than 20% of the best benchmark submitted, the bid is liable to be technically disqualified. The decision of IISER-Pune will be final and binding in this regard.