

## Setting up a high-performance cluster for scientific computing

### **Manufacturer Eligibility:**

1. Quoted OEM should have a minimum of 10 HPC installations in Top500 list from [top500.org](http://top500.org) (necessary proof of document to be enclosed with tender).
2. OEM should have a minimum 3 Installations listed in Indian Supercomputer list (<http://topsc.cdac.in/>).
3. OEM must have presence for at least 10 years in this business (proof of selling 3 HPC facilities in the past 10 years should be attached, out of these at least one sale should be older than or equal to 10 years).

### **Bidder Eligibility:**

1. Bidder should have a minimum 10 years of experience in HPC business (proof of 3 Purchase Orders along with successful installation report must be submitted, out of these at least one sale should be older than or equal to 10 years).
2. Bidder should have registered support office in India with dedicated HPC engineer capable of handling all HPC relevant issues. Necessary proof of employment should be submitted along with certification if any.
3. Bidder should be an authorized dealer or business partner of OEM.
4. Bidder should have installed at least 5 HPC systems (at least 5 compute nodes per system) in research institutes or Centers of Government of India including at least 1 system with 10 compute nodes or more.
5. Bidder must provide the estimation of the cluster performance based on the application performance using Linpack and SPEC.

### **General Terms:**

1. The cost must include all the components of the whole system. Each bid should specify the details of the configuration and should specify the cost of each component.
2. The price should also be mentioned on per node (CPU/GPU) basis. Quantity of nodes may increase or decrease at the discretion of IIT Madras.
3. The quotation should meet all the above-mentioned technical specifications. A detailed compliance sheet in the format provided has to be submitted in accordance with the above specifications. Any deviations have to be highlighted and details must be mentioned.
4. Details of power and cooling requirements for the proposed cluster should be submitted along with quote. Power socket details should be provided.

5. Complete installation and commission of HPC cluster system to be done by the bidder.

**Master Node specifications:** see specifications in the following pages

**Compute Nodes' specifications:** see specifications in the following pages

**Storage Nodes' specifications:** see specifications in the following pages

S. No.	Item	Technical Specification
1	Total Number of nodes	<p><b>1 Master Node for login, job launch and other system administration activities</b></p> <p><b>12 CPU-only compute nodes</b></p> <p><b>1 Storage node (60 TB NAS)</b></p>
2	Master node (Login, IO, etc)	<ul style="list-style-type: none"> <li>• Processors - <b>Intel Xeon Gold 6XXX series or above</b></li> <li>• Clock at least 32 instructions per core per cycle</li> <li>• At least <b>20 cores</b> per socket with 2.4 GHz or more base processor frequency</li> <li>• Two processors per node</li> <li>• At least <b>6 GB DDR4 RAM</b> for main memory in a balanced configuration per <b>core</b> with 2666 MHz or more (If such a configuration is not feasible, then at least <b>192 GB DDR4 RAM</b> for main memory in a balanced configuration per <b>node</b> with 2666 MHz or more)</li> <li>• At least 1 TB 15K RPM SAS Hard disk per node</li> <li>• Rack mountable with suitable mounting kit</li> <li>• Service level: 24x7</li> </ul>

3	CPU-only compute nodes	<ul style="list-style-type: none"> <li>• Processors - <b>Intel Xeon Gold 6XXX series or above</b></li> <li>• Clock at least 32 instructions per core per cycle</li> <li>• At least <b>20 cores</b> per socket with 2.4 GHz or more base processor frequency</li> <li>• Two processors per node</li> <li>• At least <b>6 GB DDR4 RAM</b> for main memory in a balanced configuration per <b>core</b> with 2666 MHz or more (If such a configuration is not feasible, then at least 192 GB DDR4 <b>RAM</b> for main memory in a balanced configuration per <b>node</b> with 2666 MHz or more)</li> <li>• At least 1 TB 15K RPM SAS Hard disk per node</li> <li>• One free PCI-e slot for future expansion</li> <li>• Rack mountable with suitable mounting kit</li> <li>• Redundant power supplies for all nodes</li> <li>• Service level: Next business day</li> </ul>
4	Total compute power	<p>- <b>The overall system must support expansion to at least 22 CPU-only compute nodes and 2 GPU nodes</b></p> <p>- For additional expansion in future, the vendor must offer at the bid price or below</p> <p>- Bidder must also include the total computing power in TFlops for the overall system.</p>
5	Storage Node	<ul style="list-style-type: none"> <li>- At least 60 TB NAS in total for usable file system provided in the storage node</li> <li>- At least RAID 5 across the total storage.</li> <li>- It would be split as 60% of scratch area and 40% of home storage both available in all nodes with an I/O bandwidth of 25 GBPS</li> <li>- Backup up of 40% of home with a NAS Storage of 10TB; with Necessary auto backup/restore control</li> <li>- processor: <b>Intel Xeon Silver 4XXX series or above</b> with at least 10 cores per socket and at least 4 GB RAM for main memory in a balanced configuration per core with 2666 MHz or more</li> <li>- Two processors per node</li> </ul>

6	Interconnect	<ul style="list-style-type: none"> <li>• Intel or Mellanox Infiniband switch with at least 100 Gbps bandwidth that can support future extension to double the size of the present cluster and adequate number of IB adapters for the cluster quoted.</li> <li>• Adequate redundancy to avoid single-point failure in the interconnect should be provided. All switches should have redundant power supplies</li> <li>• At least 1 G network card</li> </ul>
7	Power	<ul style="list-style-type: none"> <li>• As required for the solution above; the vendor is expected to provide the estimate of required power and cooling</li> </ul>
8	Rack mounting kit and enclosure	<ul style="list-style-type: none"> <li>• As required for the solution above and should support expansion to at least 24 nodes (22 CPU-only nodes and 2 GPU nodes). The rack should be 42U or smaller.</li> </ul>
9	Operating system/ Cluster Management	<ul style="list-style-type: none"> <li>• Licensed OS (64-bit Linux or Unix variant) and licensed cluster management/monitoring software should be fully supported (by the OEM), with updates during the warranty period. Cluster management software should be licensed without the limits on the number of nodes.</li> </ul>
10	Compilers libraries and tools	<ul style="list-style-type: none"> <li>• Support for the complete software suite including all the software currently supported on the Aqua cluster at IITM (OpenFOAM, Comsol, Mathematica, Gaussian, Ansys/fluent, Namd, Lammmps, Gromacs, Amber, Accelerys, Matlab, VASP, Quantum Espresso, CPMD, CP2K) along with various compilers (such as GNU GCC collection, intel, java compilers).</li> <li>• Also, GPU Computing and AI related software/SDK such as CUDA V9.x, 10.x, 11.x, Tensorflow/Pytorch/FastAi/XGBoost, CUDNN v7.x.</li> <li>• Modules environment package for managing software environments.</li> </ul>
11	System administration	<ul style="list-style-type: none"> <li>• All HPC components should be administered from a single management console/master node.</li> </ul>
12	Software reliability	<ul style="list-style-type: none"> <li>• Software failover for critical system services including system database, system logger and batch file systems, parallel file systems;</li> </ul>
13	Hardware reliability	<ul style="list-style-type: none"> <li>• Redundant paths for all system RAID</li> <li>• Redundant power supplies and voltage regulator modules</li> </ul>

14	Job management	<ul style="list-style-type: none"> <li>• Licensed and commercial-grade or open-source HPC workload and resource management with support for the entire lifetime of the cluster.</li> <li>• Accounting capability to summarize the total CPU usage of a user in a specific period.</li> </ul>
15	Safety	<ul style="list-style-type: none"> <li>• International standard safety level</li> </ul>
16	Onsite warranty	<ul style="list-style-type: none"> <li>• Entire HPC solution including each software, firmware and hardware component should have at least a <b>5 years warranty</b> from the HPC solution provider, from the date of acceptance.</li> </ul>
17	Installation & Support	<ul style="list-style-type: none"> <li>• The vendor should install and test following as well as any other required libraries/packages for running cluster. Operating system Cluster management tools (Open Source) Job scheduler installation C, C++, FORTRAN compilers Parallelization softwares like MPI, OpenMP and Pthread Web based management software for monitoring, managing and configuring servers.</li> <li>• The hardware of the cluster must be factory integrated, installation of the hardware, OS, software components must be done by trained engineers.</li> <li>• During the 5 year period of warranty, bidder must provide complete management support for the cluster i.e., comprehensive fault management and performance management including software required to perform these tasks.</li> <li>• The specifics of the 24 x 7 support in terms of response and resolution time for various types of problems for the HPC cluster must be mentioned separately.</li> </ul>
18	Cluster performance	<ul style="list-style-type: none"> <li>• The vendor must demonstrate the promised performance of the installed cluster based on the application performance using Linpack and SPEC.</li> </ul>