

Purchase of 10 node High Performance Computing Cluster

Manufacture Eligibility:

1. Quoted OEM should have a minimum of 10 HPC Installation in Top500 list from top500.org (necessary proof of document to be enclosed with tender).
2. OEM should have minimum 3 Installation listed in Indian Supercomputer list (<http://topsc.cdac.in/>)
3. Only OEMs should quote. If an OEM aligns with a business partner, then it should be one business partner per OEM (Authorization letter from OEM should be submitted in this regard)

Bidder Eligibility:

1. Bidder should have experience in deployment of HPC cluster in Research institutes or Centres of Govt. of India. The bidder alone or along with his OEM should have installed at least five (5) computing clusters in India in last five years. Of these at least one HPC cluster should be more than ten (10) TFlops including only CPU. A document with the TFlop calculation should be provided. A minimum of 5 copies of purchase order for HPC (within last five Years only) should be submitted followed by successful Installation Reports/Certificates along-with technical offer.
2. The Bidder should have registered support office with dedicated HPC Engineer capable of handling all HPC relevant issues. Necessary proof of employment should be submitted along with Certification if any.

General terms:

1. The **total cost must be specified including 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 specified technical specifications. A detailed compliance sheet in the format provided has to be submitted in accordance with the above specifications. Any deviations has to be highlighted and details must be mentioned.
4. Details of power and cooling requirements for the proposed cluster should be submitted along with the Technical bid. IITM will be providing the necessary Power and cooling infrastructure. 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

S. No.	Item	Technical Specification
1	Total Number of nodes	<p>1 Master Node for login, job launch and other system administration activities</p> <p>10 compute nodes 1 Storage node (100TB NAS)</p> <p>At least 30% of the compute nodes to be CPU-GPU nodes.</p>
2	Master node	<ul style="list-style-type: none"> • Processors - Intel Xeon Gold 5XXX series or above • At least 192 GB DDR4 RAM for main memory in a balanced configuration per node with 2666 MHz or more. • Two processors per node • At least 1 TB 15K RPM SAS Hard disk per node • Provisions for future expansion of the cluster to double the capacity (20 nodes) • Rack mountable with suitable mounting kit • Next business day support
2	CPU-only compute nodes	<ul style="list-style-type: none"> • At Most 7 number of CPU only compute nodes • Processors - Intel Xeon Gold 5XXX series or above • Clock at least 32 instructions per core per cycle • At least 18 cores per socket with 2.4GHz or more base processor frequency • At least 192 GB DDR4 RAM for main memory in a balanced configuration per node with 2666 MHz or more. • Two processors per node • At least 2 TB 15K RPM SAS Hard disk per node • One free PCI-e slot for future expansion • Rack mountable with suitable mounting kit • Redundant power supplies for all nodes • Next business day support

3	CPU-GPU compute nodes	<ul style="list-style-type: none"> • Either NVIDIA Tesla V100 32 GB cards (2 per GPU node) or NVIDIA A100 40/80GB cards (1 per GPU node). • At least 30% of the compute nodes (> 3 out of 10) to be CPU-GPU nodes (4 preferable configuration options listed at the end of this document) • If configuring 1 GPU card per node, then there should be a free PCIe slot in the node for further expansion later on. • At least 32 instructions per core per cycle • At least 192 GB main memory in a balanced configuration per node with 2666 MHz or more • At least 18 cores per socket with 2.4GHz or more base processor frequency • At least 2 TB SAS 15K RPM Hard disk per node • Redundant power supplies • Rack mountable with suitable mounting kit • Vendor should quote for PCI-e and NVLink separately, if applicable • Next business day support
4	Total compute power	<p>- The overall system must support expansion to at least 20 compute nodes (in the same 30-70 ratio of GPU-CPU nodes).</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 performance in TFlops for the overall system on SPEC benchmark and total GPU-CPU performance must be at least 10T Flops.</p>
5	Storage Node	<p>- At least 100TB NAS in total for usable file system provided in the storage node</p> <p>- At least RAID 5 across the total storage.</p> <p>- 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</p> <p>- Backup up of 40% of home with a NAS Storage of 10TB; with Necessary auto backup/restore control</p> <p>- Processor with at least 16 cores, at least 96GB RAM in a balanced configuration</p>
6	Interconnect	<ul style="list-style-type: none"> • 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. • Adequate redundancy to avoid single-point failure in the interconnect should be provided. All switches should have redundant power supplies

		<ul style="list-style-type: none"> At least 1 Gbps network card
7	Power	As required for the solution above; the vendor is expected to visit the site to understand the existing infrastructure.
8	Rack mounting kit and enclosure	As required for the solution above and should support expansion to at least 20 nodes. The rack should be 42U or smaller.
9	Operating system/ Cluster Management	<ul style="list-style-type: none"> 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.
10	Compilers , libraries and tools	<ul style="list-style-type: none"> Support for the complete software suite including all the software currently supported on the Aqua cluster at IITM (OpenFOAM, Abaqus, Comsol, Mathematica, Gaussian, Nastran, Dytran, Marc, Ansys/fluent, Namd, Lammmps, Gromacs, Amber, Accelerlys, Matlab.) along with various compilers (such as GNU GCC collection, intel, java compilers). Also, GPU Computing and AI related software/SDK such as - NVIDIA HPC SDK v20.x, CUDA V9.x, 10.x, 11.x, Tensorflow/Pytorch/FastAi/XGBoost, CUDNN v7.x. Modules environment package for managing software environments.
11	System administration	<ul style="list-style-type: none"> All HPC components should be administered from a single management console/master node.
12	Software reliability	<ul style="list-style-type: none"> Software failover for critical system services including system database, system logger and batch file systems, parallel file systems;
13	Hardware reliability	<ul style="list-style-type: none"> Redundant paths for all system RAID Redundant power supplies and voltage regulator modules
14	Job management	<ul style="list-style-type: none"> Licensed and commercial-grade or open-source HPC workload and resource management with support for the entire lifetime of the cluster. Accounting capability to summarize the total CPU usage of a user in a specific period.
15	Safety	<ul style="list-style-type: none"> International standard safety level

16	Onsite warranty	<ul style="list-style-type: none"> Entire HPC solution including each software, firmware and hardware component should have at least a 5 years warranty from the HPC solution provider, from the date of acceptance.
17	Installation & Support	<ul style="list-style-type: none"> The vendor should install and test following as well as any other required libraries/packages for running cluster. <ul style="list-style-type: none"> 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. The hardware of the cluster must be factory integrated, installation of the hardware, OS, software components must be done by trained engineers. 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. 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. Committed performance benchmark numbers to be demonstrated post installation as an acceptance criteria.

Options that the vendors should quote for:

- All vendors are required to bid for all the options below failing which the bids submitted will be totally rejected.
- IITM reserves all rights to choose any one of the options from the specifications given above.
- The orders will be placed only for one option and the decision on the option to be picked up will be made during the final comparison of financial bids.

Option	CPU Specs	#CPU Nodes	#GPU Nodes	GPU and #GPU cards per GPU
A	At least 18 cores per socket	7	3	NVIDIA V100 32GB (2 per node)
B		7	3	NVIDIA A100 40/80GB (1 per node)

C		6	4	NVIDIA V100 32GB (2 per node)
D		6	4	NVIDIA A100 40/80GB (1 per node)