



Indian Institute of Science Bangalore

Prof. Mayank Shrivastava
Professor
Department of Electronic Systems Engineering
Indian Institute of Science Bangalore
560012, Bangalore, Karnataka, India

Inquiry Number: DESE/LU/MSA/10/2024-25

Dated: 11.11.2024

Domestic Tender: Request for Quote from Indian Original Equipment Manufacturer (OEM) or their Indian authorized distributor for the procurement of High-Performance Computing (HPC) 1024 Core CPU Cluster with GPU Nodes

**Indian Institute of Science, Bangalore
(Last Date: 02/12/2024)**

Dear Sir/Madam,

Kindly send your best price quotation (in INR only) for the following item with various accessories on FOR-IISc Bangalore basis to the undersigned. Your quotation should clearly indicate the terms of delivery, delivery schedule, entry tax, payment terms, etc.

Your quote should also include mode of payment and **should reach the undersigned, duly signed on or before 14:00 hours (IST) on 02/12/2024.**

The quote must include all details of technical specifications of the equipment along with the commercial terms and conditions, the bill of materials, printed technical brochure and any other supporting document. **Please enclose a compliance certificate, printed on your letter head, along with the quote.**

The commercial bid must include the price of the item in Indian currency, indicating the following separately:

- FOR price
- Insurance
- Post warranty maintenance charges
- Tax
- Total

The quotation should address to:

The Chairman,
Department of Electronic Systems Engineering
Indian Institute of Science, Bangalore – 560012



Indian Institute of Science Bangalore

Prof. Mayank Shrivastava
Professor
Department of Electronic Systems Engineering
Indian Institute of Science Bangalore
560012, Bangalore, Karnataka, India

1. Introduction and Scope

- **Purpose:** This tender is issued for the procurement, installation, and commissioning of a High-Performance Computing (HPC) Cluster consisting of 1024 CPU cores and dual GPU nodes. The system will support intensive computational tasks in research and data processing.
- **Scope of Work:** Supply, installation, testing, commissioning, and demonstration of a fully operational HPC cluster. The vendor shall also provide training and support for the system.

2. Technical Specifications

A. Hardware Specifications

1. Master Node (1 Node)

- **Processor:** 2 x Intel Xeon Gold Gen3-6338 (32 cores, 64 threads, 2.0 GHz, 48M cache)
- **Memory:** 256GB DDR4 ECC, 3200 MHz in homogenous fully balanced mode
- **Storage:**
 - Primary: 960GB m.2 NVMe
 - Additional: 6 x 8TB Enterprise SATA HDDs @7200RPM
 - Require Raid Controller Card for these HDD's
- **Networking:**
 - 1 x Mellanox 100Gbps EDR adapter
 - 2 x 1GbE ports
- **Power Supply:** 1 x 800W RPS (80 PLUS Platinum in (1+1) mode)
- **Chassis:** 2U Rack-mount with rail kit

2. Compute Nodes (16 Nodes)

- **Processor:** 2 x Intel Xeon Gold Gen3-6338 per node
- **Memory:** 256GB DDR4 ECC, 3200 MHz in homogenous fully balanced mode
- **Storage:** 960GB m.2 NVMe SSD
- **Networking:** Mellanox 100Gbps EDR adapter, and 2 x 1GbE ports
- **Power Supply:** 1 x 800W RPS (80 PLUS Platinum in (1+1) mode)
- **Chassis:** 2U Rack-mount with rail kit



Indian Institute of Science Bangalore

Prof. Mayank Shrivastava
Professor
Department of Electronic Systems Engineering
Indian Institute of Science Bangalore
560012, Bangalore, Karnataka, India

3. GPU Nodes (2 Nodes, 8 GPUs Each)

- **Processor:** 2 x Intel Xeon Gold Gen3-6338 per GPU node
- **Memory:** 256GB DDR4 ECC, 3200 MHz in homogenous fully balanced mode
- **Storage:**
 - Primary: 960GB m.2 NVMe
 - Additional: 4 x 8TB Enterprise SATA HDDs @7200RPM
 - Require Raid Controller Card for these HDD's
- **GPU:** 8 x NVIDIA A6000-48GB ADA per node (Unit price of Nvidia 6000 ADA should be mandatorily given separately for arriving at exact number of GPUs to buy)
- **Networking:** Mellanox 100Gbps EDR adapter, and 2 x 1GbE ports
- **Power Supply:** 4 x 2000W RPS (80 PLUS Platinum in (3+1) mode)
- **Chassis:** 4U Rack-mount with rail kit

4. Interconnect

- **Switch:** Mellanox 36-port 100Gbps EDR InfiniBand switch. This should be rack mountable.

5. NAS Node: Provide Optional Price of Adding a rack mountable NAS Node (Please note that this is an optional item) with a Storage of 14 x 10TB Enterprise-grade SAS HDDs, 100TB usable capacity in RAID-6 configuration

6. Rack Enclosure: Provide Optional Price of (2 units) 42U 1000mmx 600mm server with minimum 40 x c14 socket vertical PDUs. Front and rear doors of the rack should be perforated and the sides louvered.

All the 5 components above (master, compute, GPU nodes, switch and NAS node) should be mountable into an existing rack of dimension 600 X 1200 mm, which has 32A, 3 Phase, high density PDU x 02 nos with each PDU consisting of 18 Nos of C-13 sockets and 6 nos of C-19 sockets.

B. Software Requirements

Operating System

- Options: CentOS 8 (or CentOS Stream 8), Rocky Linux 8 (recommended for stability in scientific computing).
- Purpose: Provides a stable and compatible environment for high-performance computing (HPC) applications, including cluster management tools.

Cluster Management and Job Scheduling Software

E-mail : mayank@iisc.ac.in (write to msdlab.ese@iisc.ac.in for tender related queries)
Phone : +91-80-2293-2732
Faculty Web : <https://faculty.dese.iisc.ac.in/mayank/>
Institute Web : <http://www.iisc.ac.in/>



Indian Institute of Science Bangalore

Prof. Mayank Shrivastava
Professor
Department of Electronic Systems Engineering
Indian Institute of Science Bangalore
560012, Bangalore, Karnataka, India

- OpenHPC: A comprehensive open-source cluster management toolkit that includes essential HPC components.
- SLUR/OpenPBS/SGE: Any open source Job submission engine
- Ganglia: For cluster monitoring, Ganglia provides real-time monitoring and performance visualization, which is critical for managing a large cluster.

Networking and Communication Libraries

- MPI (Message Passing Interface):
 - OpenMPI or MPICH: Both are widely used open-source implementations of MPI. MPI is crucial for enabling inter-process communication across nodes, especially in parallel computing workloads.
- OFED (OpenFabrics Enterprise Distribution): Provides optimized InfiniBand drivers and libraries for high-speed networking. Essential for configuring the Mellanox 100Gbps EDR InfiniBand network.

Storage and File Systems

- NFS (Network File System): For simpler file sharing, NFS can be used to mount shared directories across nodes, though it does not provide the same performance as parallel file systems.

GPU Software and Libraries

- NVIDIA CUDA Toolkit: Essential for running GPU-accelerated applications on the A6000 ADA GPUs. CUDA provides libraries and tools for deep learning, molecular dynamics, and other GPU-intensive applications.
- cuDNN: NVIDIA's GPU-accelerated library for deep learning. This will be useful for applications like M3GNet.
- NCCL (NVIDIA Collective Communications Library): Optimizes multi-GPU communication and is especially useful in multi-GPU setups like this one.

Scientific and Computational Libraries

- MKL (Math Kernel Library): Intel's optimized library for math and linear algebra functions, which accelerates scientific computations.
- OpenBLAS: An open-source alternative to MKL for optimized linear algebra routines.
- FFTW: A library for computing the discrete Fourier transform (DFT), useful in various scientific applications.
- HDF5: A file format and set of tools for managing complex data, commonly used in scientific computing for data storage.



Indian Institute of Science Bangalore

Prof. Mayank Shrivastava
Professor
Department of Electronic Systems Engineering
Indian Institute of Science Bangalore
560012, Bangalore, Karnataka, India

Additional Tools

- **Ansible:** Automates configuration and deployment across the cluster. Useful for managing and updating nodes simultaneously.
- **Git:** Version control system, essential for managing software and code repositories across the cluster.
- **Python (Anaconda Distribution):** Provides a comprehensive environment with scientific libraries (NumPy, SciPy, Pandas) that are widely used in research.
- **JupyterHub:** Allows multiple users to run Jupyter notebooks in a shared environment, which can be helpful for research and analysis tasks on the cluster.

3. Installation, Commissioning, and Testing

- **Installation:** Vendor must perform racking, stacking, cabling, software installation, and full system integration. All cables for interconnecting the systems to cluster should be considered in the quote.
- **Commissioning:** Complete system setup, configuration, and connectivity testing to ensure full functionality.
- **Performance Testing:**
 - Benchmark tests on CPU, GPU, and network interconnect performance.
 - Application-level performance tests for scientific libraries, MPI, and machine learning frameworks.
- **Acceptance Criteria:** System must pass all specified performance benchmarks and demonstrate stability and efficiency in running target applications.
- **No additional cost for S/W & H/W Installation, Commissioning and Performance Testing**

4. Warranty and Support

- **Warranty:** Minimum 3 years onsite, next-business-day support on all components.
- **Extended Maintenance:** Annual Maintenance Contract (AMC) option for two additional years post-warranty. Quote this as an optional items.
- **Direct OEM Support:** OEM must provide direct support in case of any failure in service from the vendor.

5. Commercial Terms and Conditions

- **Pricing:** Provide detailed, itemized pricing, including per-node, per-GPU, and other specific components, wherever required.



Indian Institute of Science Bangalore

Prof. Mayank Shrivastava
Professor
Department of Electronic Systems Engineering
Indian Institute of Science Bangalore
560012, Bangalore, Karnataka, India

- **Detailed Specs:** Include detailed specifications for each hardware component (Master node, Compute nodes, GPU nodes, InfiniBand switch)
- **Bid Evaluation:** Evaluation will be based on both technical compliance and commercial viability.
- **Scalability Clause:** The option to add additional compute nodes or GPUs and at the same unit price.
- **Payment Terms:** 90% payment upon delivery and 10% upon successful installation, testing, and acceptance.
- **Delivery Timeline:** 4 weeks for delivery and an additional week for complete installation.
- **Penalty for Delays:** A penalty will apply for each day beyond the specified delivery timeline. This procurement is from a CSR fund and hence it's closure before March 1st 2025 is mandatory.

6. Eligibility Criteria

- **Experience:** The vendor must have prior experience with at least 3 HPC installations, with at least one cluster of 256 cores or more.
- **Financial Requirements:** Minimum annual turnover of ₹5 Crores for the past three financial years.
- **Local Presence:** Must have a local office and support team in Bengaluru.
- **Bidder Authorization:** Only bids from OEMs or authorized Indian distributors will be accepted. MSE OEM bidder will be given preference.

8. Documentation and Compliance

- **Bids:** The tender is two-bid tender with Technical and Commercial bids to be submitted in two separate envelopes. The commercial bids of technically disqualified bidder will not be opened.
- **Compliance Certificate:** Signed component level compliance matrix covering all technical specifications should be submitted in the technical bid.
- **Reference Sites:** Provide references for three cluster installations completed in the past three years.
- **OEM Authorization:** Bids from OEM authorized bidders will only be entertained. Manufacturer Authorization Form should be enclosed.

9. Important Notes

- **Submission Guidelines:** Submit technical and commercial bids in separate sealed envelopes.
- **Confidentiality:** Ensure confidentiality of all documents and information provided.
- **Rights Reserved:** The issuing authority reserves the right to modify the specifications or quantity requirements.



Indian Institute of Science Bangalore

Prof. Mayank Shrivastava
Professor
Department of Electronic Systems Engineering
Indian Institute of Science Bangalore
560012, Bangalore, Karnataka, India

PI Terms and conditions (should be included in compliance certificate):

1. To an extent possible, share itemized pricing for every item (this will allow us to add additional resources in the PO, such as RAM, HDD, etc., if the budget allows, before placing the order).
2. Necessary training to operate the procured setup should be provided without additional cost.
3. In principle onsite installation should be free of cost. The amount of time / day committed by the engineer during installation must be clearly stated.
4. Software upgrade, if any, must be free of cost for next 3 years.
5. The vendor must assure that there are no bugs and glitches with the integration. In case of glitches or bugs at the time of installation, vendor must fix the issues in less than three days from the start date.
6. In case of hardware/software issues or support, vendor should be able to provide required solution within three days.
7. Additional quote for an annual maintenance contract beyond the warranty period should be included for the next 5 years.
8. The vendor should have a good track record of delivering such equipment at universities/research institutions (please furnish the details).
9. Please provide list of customers who have procured your equipment in last 5 years.
10. The vendor should be able to repair and maintain the equipment, once it is installed. No travel claims must be made by vendor for servicing during the warranty/guarantee period.
11. The system must be delivered at the earliest. The smallest lead time will be appreciated. Our expectation is shipment within 4 weeks after PO and full or part payment post installation.
12. On all systems the payment terms will be specified in the commercial proposal and is subject to negotiation.
13. The validity period of the quotation should be 90 days at least.
14. Please provide details of the number of trained personnel, who can service the setup.
15. See other Terms & Conditions, guidelines, eligibility criteria etc. in enclosed document in the next pages.



Indian Institute of Science Bangalore

Prof. Mayank Shrivastava
Professor
Department of Electronic Systems Engineering
Indian Institute of Science Bangalore
560012, Bangalore, Karnataka, India

IISc Terms and conditions:

1. The Bidder should belong to either Class-1 or Class-2 suppliers distinguished by their “local content” as defined by recent edits to GFR. They should mention clearly which class they belong to in the cover letter. a) Class-1 supplier: Goods and services should have local content of equal to or more than 50%. b) Class-2 supplier: Goods and services should have local content of equal to or more than 20 % and less than 50%.
2. Bidders offering imported products will fall under the category of non-local suppliers. They cannot claim themselves as Class-1 local suppliers/Class-2 local suppliers by claiming the services such as transportation, insurance, installation, commissioning, training, and other sales service support like AMC/CMC, etc., as local value addition.
3. Purchase preference as defined by the recent edits to GFR (within the “margin of purchase preference”) will be given to the Class-1 supplier.
4. MSMEs can seek an exemption to some qualification criteria. IISc follows GFR2017 for such details.

Sincerely,

Prof. Mayank Shrivastava
Professor
Department of Electronic Systems Engineering,
Indian Institute of Science,
Bangalore, Karnataka 560012, India
Secretary (Ms. Rekha's) Contact: 9972525771
(On Behalf of Purchase Committee)
Email: msdlab.es@iisc.ac.in (for tender related queries)