



**INDIAN INSTITUTE OF TECHNOLOGY BOMBAY
MATERIALS MANAGEMENT DIVISION
Powai, Mumbai 400076.**

Reference No. 2021-22/121 PR No. 1000020460 (Rfx No. 6100001192)

Detailed Technical Specifications for High Performance Computing Cluster:

Introduction: IIT Bombay desires to augment its High Performance Computing Resources. A facility of this kind is expected to offer a boost to the modeling and simulation research efforts encompassing several departments. Some of the proposed activities include computational fluid dynamics, bio-molecular simulations, computational chemistry and biology, transportation engineering, nano-electronics and several other areas of science and technology (Espresso, VASP, LAMMPS, Siesta, CPMD, CP2K, ABNIT, NAMD, GROMACS, AMBER, Gaussian 09, Gaussian 16, ANSYS etc.). Therefore, we intend to procure maximum (sustained) compute performance within the power, cooling and space envelope mentioned in the technical specifications.

Important Infrastructure Criteria:

- A. **Cooling:** IIT Bombay will provide chilled water through a water pipe loop (of 200 mm diameter) with an inlet temperature of 15 degree and a delta of 5 degree Celsius to cool the IT components. Winning bidder has to arrange for additional plumbing (if required) to supply the chilled water to the rack.

Please note, IIT will not accept a proposition or request for conventional cold aisle – hot aisle arrangement.

- B. **Power:** Maximum available power of 350 KW (hard limit) is available for the computer hardware through a series of redundant local DB. Arrangements for additional cabling (if required) for powering server racks have to be arranged by the winning bidder.

Bidders are advised to visit the data centre at a mutually convenient time, **once**, before submitting the bid to clarify some of these points. No photograph or video graphing is allowed during such a site visit.

I. Technical Specifications

No	System Attribute		Specification
1.		Processor	At least 2 tera-FLOPs (theoretical peak) per socket per node. Processor has to be PCIe gen 4.0 compliant
		Memory	Minimum 256 GB per node 3200 or more MHz RDIMM (multi rank) and protected by advanced ECC, Chip Kill or equivalent technology. The memory DIMMs must be installed in a fully balanced mode to operate at the maximum rated frequency of the memory. Also, all the available memory channels should be populated with identical capacity DIMMs
	(A) CPU-only Compute Nodes	No of Processors per Node	One or two processors per node
		Interconnect	As required for the optimized solution
		Type	Rack mountable/Blade enclosure/Dense Server Chassis with suitable mounting kit
		Power Supply	Redundant hot swappable power supply and redundant fans and appropriate number of properly configured power distribution units (PDUs) for the enclosures/racks solution
		Warranty and Service Level	Warranty: As described in SLA. Service Level: 24 x 7; Complaints have to be attended within 4 hours of receiving. Rest are described in SLA.
	(B)	Processor	At least 2 teraFLOPs (theoretical peak) per socket. One or two processors per node. Processor has to be PCIe gen 4.0 compliant.

	CPU+Accelerator Compute Nodes	Co-Processor/ Accelerator	Performance from each accelerator should 9TF (FP64 cores) or higher Minimum 64GB (HBM2 or better) per accelerator or higher Each accelerator node should be configured with 4 accelerators per node or higher.
		Memory	Minimum 512 GB per node 3200 or more MHz RDIMM (multi rank) and protected by advanced ECC, ChipKill or equivalent technology. Also should be scalable up to 1024 GB per node. The memory DIMMs must be installed in a fully balanced mode to operate at the maximum rated frequency of the memory. Also, all the available memory channels should be populated. All the Memory channels should be populated with identical capacity
		Interconnect	As required for the optimized solution
		Type	Rack mountable/Blade enclosure/Dense Server Chassis with suitable mounting kit
		Power Supply	Redundant hot swappable power supply and redundant fans and appropriate number of properly configured power distribution units (PDUs) for the enclosures/racks solution
		Warranty and Service Level	Warranty: As described in SLA. Service Level: 24 x 7; Complaints have to be attended within 4 hours of receiving. Rest are described in SLA.
	(C) High memory CPU-only Compute Nodes	CPU based compute node as described in S. No. 1A with higher memory of 1024 GB per node	

<p>Total number of compute nodes</p>	<p>Minimum 2 PetaFLOP of compute power (theoretical peak) arranged in the following manner: (A) CPU only nodes, (B) hybrid CPU+Accelerator nodes and (C) minimum of 10 (ten) high memory CPU nodes is expected to be purchased. 20±1% of compute power (as measured in rated TeraFlopsS) should be derived from Accelerators alone and 80±1% from CPUs of CPU-only nodes (including High Memory nodes). The overall memory in the solution from Compute Nodes should be at least 125 TB (1TB = 1024GB) and a minimum of 25000 cores.</p>		
<p>2.</p>	<p>Service Nodes (Master, login, I/O etc)</p>	<p>Processor</p>	<p>As required for optimized solution</p>
		<p>Memory</p>	<p>Minimum 512 GB per node 3200 or more MHz RDIMM (multi rank) and protected by advanced ECC, ChipKill or equivalent technology. Also should be scalable up to 1024 GB per node. .</p> <p>The memory DIMMs must be installed in a fully balanced mode to operate at the maximum rated frequency of the memory. Also, all the available memory channels should be populated. All the Memory channels should be populated with identical capacity</p>
		<p>No of Nodes</p>	<p>Minimum ten (10).Higher as may be required for the solution with optimum parallel file system (i.e. master node(s), login node(s), I/O node(s), cluster management nodes(s), etc.) with adequate redundancy.</p>
		<p>Type</p>	<p>Rack mountable/Blade enclosure/Dense Server Chassis with suitable mounting kit.</p>
		<p>Power Supply</p>	<p>Redundant hot swappable power supply and redundant fans and appropriate number of properly configured power distribution units (PDUs) for the enclosures/racks solution</p>
		<p>Warranty and Service Level</p>	<p>Warranty: As described in SLA.</p> <p>Service Level: 24 x 7; Complaints have to be attended within 4 hours of receiving. Rest are described in SLA.</p>

<p>3.</p> <p>Storage</p>	<p>Overall 5 petabyte (PiB) of usable file system to be provided for the HPC cluster with SAS/NL-SAS. The overall storage system should be able to deliver 100 Gbps of write throughput (in 1 MB block size) across 5 PB proposed storage systems.—It should support hardware/software RAID with at least RAID6 across 5 PiB in 8D+2P array or equivalent. Redundant power, cooling and network ports required. Separate management switches for storage are required. Storage has to be housed in separate racks.</p> <p>3PB of <i>scratch</i> area using a parallel file system (Commercially supported Lustre or Spectrum Scale only) over hot-swap SAS/NL-SAS disk and connected via suitable interconnects. File system storage should be connected to master/login and all compute nodes directly through a high bandwidth, low latency network of switches only To be exported as PFS volume to all the compute nodes through suitable interconnects and as NFS volume to the non-compute nodes.</p> <p>2PB of usable <i>home</i> storage to be implemented using cluster/parallel file system (Commercially supported Lustre or Spectrum Scale only) over hot-swap SAS/NL-SAS disk. To be exported as PFS volume to all the compute nodes through suitable interconnects and as NFS volume to the non-compute nodes.</p> <p>The configuration of the storage solution should be in “no single point of failure” mode.</p> <p>Storage management should be part of the proposed solution.</p> <p>2% of total capacity needs to be additionally provided as Global Hot spare / Hot Space.</p> <p>The following are required for Lustre PFS. Equivalent terminology may be adopted for PFS other than Lustre. OSS and MDS nodes be configured as per architectural requirement of proposed parallel File System other then Luster</p> <ol style="list-style-type: none"> 1. MDT should be mounted only with the MDS server. 2. OST should be mounted only with OSS servers. 3. For MDT Fail over, MDS Nodes should be configured with an active/passive pair. Metadata should be configured on SSDs only in a separate Storage Array. Metadata and OST related data should not use a common Storage Array 4. For OST Fail over, OSS Nodes should be configured with an active/active pair. 5. If 3 and 4 are provided, they should also satisfy the following conditions <ol style="list-style-type: none"> a. High Availability should be automated. b. File system should not go down, even if one or more IO nodes fail. c. Mounting and unmounting of the file system should have happened without error.
---------------------------------	--

		<p>Warranty and Service Level:</p> <p>Warranty: As described in SLA.</p> <p>Service Level: 24 x 7; Complaints have to be attended within 4 hours of receiving. Rest are described in SLA.</p>
4.	Interconnect	<p>Primary Communication Network (In fat-tree topology or any other proprietary well established proventopology connecting all Compute Nodes, Service Nodes & I/O Nodes with a minimum bandwidth of 200 Gbps per port) [Bidder should describe the design of interconnects in detail]</p> <ol style="list-style-type: none"> 1. Adaptive routing in the interconnect must be comprehensively supported. It must be turned on while running application benchmarks. 2. Cables: Required length and number for the proposed solution 3. Required number of power cables of IEC C13 to C14 type cable with proper power rating 4. Required rack mounting kit <p>Cluster Management Network</p> <ol style="list-style-type: none"> 1. Switch(es) with required number of ports 2. Cables: Required length and number for the proposed solution 3. Required number of power cables of IEC C13 to C14 type cable with proper power rating 4. Required rack mounting kit <p>Adequate redundancy to avoid any single-point failure in the interconnect, should be provided. All Network/Interconnect cabling must be structured and adhere to ANSI/TIA – 568 standard.</p>
5.	Operating System	<p>64 bit Commercially supported Linux or Unix variant with support along with software for monitoring and maintenance of the entire system, providing single system view and fault monitoring.</p>
6.	Compilers, Libraries and Tools	<p>Complete software suite including support for FORTRAN 77, 90, 95, 2003 including Co-array FORTRAN described in the FORTRAN 2008 specification; licensed C/C++, must support Unified Parallel C and other standard MPI libraries. Separate CPU and Accelerator based compilers should be provided. Minimum 5 user commercial licenses for compilers.</p>

7.	System Administration	<p>System administration, management and monitoring of all HPC components from single management node/console(s) (both CLI and GUI based).</p> <p>Hardware Supervisory System with independent management fabric between all system blades and cabinet level controllers (both CLI and GUI based).</p>
8.	Software Reliability	<p>Node Knowledge and Reconfiguration facilities.</p> <p>Should have provision to monitor operating system kernels from one central system.</p> <p>Software fail over for critical system services including system database, system logger and batch subsystems, parallel file system object storage fail over, file metadata server fail over</p>
9.	Hardware Reliability	<p>Redundant paths to all system RAID</p> <p>Redundant power supplies; redundant voltage regulator modules</p>
10.	Job Management	<p>Licensed and Commercial-grade HPC workload and resource management. The job scheduler must allow multiple jobs of the same user or different users on a given compute node (CPU only as well as CPU-Accelerator).</p>
11.	Power	<p>350 KW (hard limit) AVAILABLE for entire solution, But lesser is better (for Compute & Storage)</p>
12.	Cooling	<p>100 TR. of cooling through Chilled water supply to rack at 14 degree C</p>
13.	Safety level	<p>FCC Class A, VCCI Class A, ICES-003, EN 50022:2006 Class A, AS/NZS CISPR 22:2006, EN 55024:1998 +A1:2002 +A2:2003 (whichever is applicable) or equivalent international standard safety level.</p>

<p>14.</p>	<p>Installation and acceptance</p>	<ol style="list-style-type: none"> 1. Installation by trained brand engineers of the HPC Solution provider (henceforth termed as OEM). 2. Demonstration of the full capabilities of the system that are listed in the proposal by the OEM brand engineers. 3. Installation of OS, Job scheduler and other management utilities should be done by the trained OEM brand engineers, prior to the delivery of the factory integrated solution. For security reasons and firewall related issues, installation and testing of scientific applications and any other application/utility, other than the factory integrated ones, should be done on site. 4. Cluster functionality tests need to be performed by the trained engineers from OEM. 5. The proposed solution should be validated and certified by OEM. 6. After deployment the same code(s) (which were given for benchmark testing) will be run on the new system in the presence of IIT Bombay representatives and the benchmark data should meet or exceed expectations as projected in the technical bid. 7. Storage throughput has to be tested on IOR/IOZone benchmark with 1MB block size of files that are greater than at least twice the total cache provided in the storage system. 8. After installation and above test runs are over, end users from IIT Bombay may check the software which we use, within a period of 7 days. 9. Training for general system administration with adequate documentation including tasks such as user/node management, installation/upgrade, queuing system management and file system management should be provided by the OEM. 10. Finally, up to 14 Days non-stop acceptance test with availability equal or greater than 99% has to be run on the delivered system. If the availability drops below 99% during the acceptance, the supplier will have to submit an official report and rectify the fault.
<p>15.</p>	<p>Reliability and Server Health Care</p>	<ol style="list-style-type: none"> 1. Dedicated, independent 1 Gbps or above system management network for nodes and cabinet level controllers in the system. Provision for monitoring racks status (i.e. cooling, power, etc) has to be there. 2. Compute nodes should be hot swappable or have provision of built-in standby nodes of the same configuration. 3. Integrated supervisory software tools to allow management, administration, monitoring and maintenance of the system at hardware and software levels. Such tools will have interfaces accessible to qualified users and allow:

		<ul style="list-style-type: none"> a. Monitoring the physical state (such as power and temperature) of each node, allowing each node to be turned on and off remotely. b. Monitoring and administering the file sharing system. c. Simultaneously installing a software package on all nodes.
<p>16.</p> <p>Additional Characteristics of the overall solution</p>	<p>All major components should provide sufficient redundancy. There should be No Single Point of Failure in the total HPC solution</p>	
	<p>System health monitoring infrastructure – Integrated supervisory system at hardware and software levels.</p>	
	<p>Flexibility in running ISV applications, legacy codes as well as highly scalable applications efficiently.</p>	
	<p>Factory Integration and testing – The total hardware solution should be factory integrated and tested before shipping to the site (even if the bidding OEM outsources any part of the solution). Certificate confirming the same must be submitted with the B.O.M.M at the time of delivery.</p>	
	<p>Programming Environment – Optimized Compilers, Tools and Libraries that improve programming productivity, promises application scalability and performance.</p>	
	<p>Firmware – All hardware should be installed with the latest stable version of firmware. Support for the firmware should be provided for the 3 yrs warranty and during the AMC period.</p>	
	<p>Upgradability – Proposed system should offer flexible upgrade options to take advantage of next generation compute processor, accelerator, I/O technology etc. Please describe possible procedures and schedules of upgrade options and the impact to the end user during such upgrades.</p>	

	<p>Warranty – The entire HPC solution including each software, firmware, and hardware component should have a 3 year warranty from the HPC Solution provider (with 24x7 service level unless specified otherwise earlier.)</p> <p>Submit the specifics of your “24x7” support in terms of response and resolution time for various types of problems in the HPC cluster.</p> <p>Additional 4th, 5th, 6th, 7th and 8th year AMC (to be quoted separately, may be purchased by IIT Bombay at the beginning of the respective year at the price quoted in this bid.)</p>
--	--

Documentation to be provided (After installation)

1. User Creation/Deletion/Modification.
2. User usage accounting – Storage and compute
3. Switching on, starting up and shutting down the cluster.
4. Disk, health status monitoring of Master/IO nodes and storage enclosure.
5. Basic troubleshooting for storage and job scheduler.
6. Step by step installation guide for node configuration from scratch.
7. Any other document/manual useful for daily administration.

II. Services Level Agreement (SLA) & Warranty

All the following conditions must be agreed upon.

1. The entire HPC solution including each software, firmware, and hardware component should have a 3 year warranty from the HPC Solution provider (with 24x7 service level unless specified otherwise earlier.)
2. The OEM (i.e. the overall HPC Solution provider) warrants that all the goods are new, unused, and of the most recent or current supported models, and that they incorporate all recent improvements in design and materials, unless provided otherwise in the Contract.
3. The OEM has to ensure that the proposed solution delivers an uptime guarantee of 98% of the entire system on a yearly basis (i.e. annual node-hours of uptime) and minimum of 95% on a monthly basis (i.e. monthly node-hours of uptime). Every percentage point of downtime between 2% to 10% on a yearly basis will incur a penalty of 0.1% of the total cost of this tender. Every percentage point of downtime above 10% on a yearly basis will incur a penalty of 0.5% of the total cost of this tender.
4. In the event of failure of any of the subsystems or components of the proposed solution, the OEM has to ensure that the defects are rectified within two full working days. Any delay in node warranty servicing beyond 3 days will incur a penalty of 0.2% of the total cost per day of delay. Any delay in storage or any of its subsystems not working beyond 24 hours will incur a penalty of 0.2% of the total cost of this tender for every completed 24 hours.
5. The defects, if any, during the guarantee/warranty period are to be rectified free of charge by arranging free replacement wherever necessary. It should be completed within 2 working days for individual servers and next working day for critical components like power supply, networking and storage after the intimation of fault.
6. During the warranty period, OEM will have to undertake comprehensive maintenance of the entire hardware components, equipment, software support and accessories supplied by the OEM at the place of installation of the equipment.
7. A letter of commitment for five years from the date of installation, with respect to Hardware Software, and Firmware support from OEM should be enclosed in the cover for Technical bid. The Offer will be rejected if the OEM fails to attach a letter of commitment.
8. The maximum penalty for non-performance will be 5% of the total cost. On reaching this limit in any year, the OEM will be considered in breach of the contract and appropriate legal action will be taken. The penalty will not apply if the delay is caused by IIT Bombay.
9. Technical support should be provided by the awarded OEM for system administration/maintenance of the HPC solution during the entire warranty period and AMC period as well.
10. The awarded OEM should protect the data stored in the system, during any upgrades of hardware/firmware/OS.
11. An inventory of common parts that require replacement shall be made available to IIT Bombay a priory on site.
12. The bidder should also provide the cost of an annual maintenance contract (AMC) within the bifurcation document requested in stage-II (Financial evaluation through

Reverse Auction)which starts at the end of the warranty period with back-to-back, onsite support from the OEM.

III. General terms and conditions:

Bidders are advised to read the following clauses carefully. Submitting your solution implies that you agree to act as per the terms and conditions mentioned below.

1. The solution has to fit into an envelope of 350 KW (hard limit) online UPS for the IT components (compute, storage, interconnect, master node), 100 TR cooling through scroll chillers operating in adequate redundancy mode and delivering chilled water to the rack at 14 degree C.
2. OEM must visit the data centre at a mutually convenient time, once. No photograph or video graphing is allowed during such a site visit.
3. Based on their visit, the OEM should submit the site compatibility proforma for the existing civil, mechanical, electrical and plumbing facilities as part of the technical bid. Any minor installation related changes that are required for the solutions should be communicated to IIT Bombay in this proforma.
4. If any minor changes are required in the data centre, it would be the sole responsibility of the OEM to carry out these changes in consultation with IIT Bombay personnel. Any such work should be finished before shipment of the hardware. The expenditure for these changes / alterations should be incurred by the OEM.
5. IIT Bombay plans to buy a complete HPC solution with a theoretical peak of 2 PetaFLOP that includes only-CPU, high memory CPU and hybrid CPU+Accelerator nodes as specified in 'Technical Specifications' in this document. The solution should produce sustainable performance not less than 75% of theoretical peak for CPU only nodes (including high memory nodes); If it produces less than 75%, additional compute nodes (CPU only) have to be supplied to compensate for the difference. Bidders should quote the total solution price for this minimum configuration. The commercial bid used in the overall techno-commercial scoring will be the bid for this minimum configuration. Bidders should also quote an item-wise unit price for all items in the solution.
6. The bidding OEM shall bear all costs during the preparation and submission of the proposal, site visit (if required) etc.
7. All the Proforma along with a covering letter should be submitted on the company letter head of the bidding OEM, along with the technical proposal. The bidding OEM has to attach verifiable documents to support their claims.
8. Bidding OEM is expected to specify all the components required to meet the objectives mentioned in the scope of the work. List all the proposed components (hardware and software) clearly as a part of their proposal. They should also describe (in brief), as far as possible, the make and model number, technical specifications etc. of these components.
9. The bidding OEM may be asked to come to IIT Bombay and present the proposed solution that they intend to supply.
10. No new information will be accepted from the bidding OEM after submission of the bids. However, IIT Bombay may ask for clarifications if required, on submitted information in order to evaluate the bid. The bidder should respond to such a clarification request within the specified time.

11. If the prices offered are not competitive or are higher than the market prices then IIT Bombay reserves the right to cancel part or whole of the tender.
12. IIT Bombay reserves all the rights to accept or reject, in full or in part, any or all the offers if
 - a. The OEM fails to comply with any material term of the contract;
 - b. The OEM fails to deliver the material(s) or any part thereof within the stipulated delivery period and /or fails to replace/ rectify any rejected or defective material(s) promptly;
 - c. The OEM fails to produce the claimed performance after installation and commissioning of the provided solution.
 - d. The OEM becomes bankrupt or goes into liquidation or the seller makes a general assignment for the benefit of the creditors or a receiver is appointed for any substantial property owned by the OEM;
 - e. The OEM has misrepresented to the IIT Bombay.
13. IIT Bombay also reserves the right to re-issue the tender without any explanation. The OEM will not have any right to object to such re-issue of tender.
14. Due to an extremely strict deadline for consuming the fund, IIT Bombay has the right to cancel the PO if the delivery, installation and acceptance testing is not completed within the stipulated timeline. Specifically
 - a. Delivery should be within 6 months of issuing of LC.
 - b. Installation, commissioning, and acceptance testing should complete within four weeks of delivery.
15. Warranty period is to be counted from the date when the installation is completed and acceptance certificate has been issued by IIT Bombay.
16. The installation has to be executed by certified and trained engineers from OEM for such HPC cluster stack and other peripherals (viz. Implementation of parallel file system, configuring primary/management network and I/O module, etc) followed by a well documented, comprehensive user training.
17. The OEM must submit the name of the engineers who are going to be deployed and will be in-charge to execute the HPC installation, along with a proof of their employment and designation, their contact details in India, working knowledge of basic HPC setup and management (viz. Job submission, Job execution, queue management, Cluster setup etc.).
18. The OEM will provide an undertaking that OEM is responsible for a 5 year performance guarantee. In case of merger/sale of business by the OEM, the above-said warranty, AMC and SLA will be applicable to the new OEM. If the new OEM does not honor the said warranty, AMC, and SLA, the IIT Bombay reserves the right to blacklist both the OEM and reserves the right to take proper legal action. The OEM has to accept this clause and submit a declaration. Failing to do the same will result in disqualification of the bid by defaulter OEM.
19. Any item not specifically mentioned in the specification but is required for successful implementation of the HPC solution (in the solution proposed by OEM) must be brought to our notice and quoted accordingly.
20. At the time of installation, if it is found that some additional hardware or software items are required to meet the operational requirement of the solution, but not included in the

OEM's original list of deliverables, the OEM shall supply such items to ensure the completeness of the configuration at no extra cost and within the installation timeline.

21. Entire installation should be done at the proposed site only. Requests for remote access for installation / fine tuning will not be entertained during the installation period.
22. Acceptance certificate(s) will be issued only after all the acceptance tests as outlined earlier (in Technical Specifications Point 15, and elsewhere in this document) have been satisfactorily passed.
23. The IIT Bombay reserves the right to publish the information about the unsatisfactory service by the bidder/OEM and action taken by the institute on their website and in the national newspaper(s).

TENDER EVALUATION:

The competent authority will evaluate all the proposals to determine whether these are complete in all respects as specified in the tender document. Evaluation of the proposal shall be done in two stages as

(a) Stage - I (Technical Evaluation):

1. Institute shall evaluate the technical bid(s) to determine whether these are meeting the essential eligibility criteria, whether any computational errors have been made, whether all the documents have been properly signed & stamped, whether all the documents as mentioned / or required to be submitted with technical bid are submitted and whether a bid is complete and generally is in order. There will not be any further technical evaluation will be done in case of incomplete bid and the bidder will be disqualified.

2. After evaluating the performance parameters offered, support structures, technical evaluation of proposed server, and references. Bidders with minimum score will be qualified to participate in Commercial bidding phase and the disqualifications will be informed to the concerned bidders.

(b) Stage- II (Financial evaluation through Reverse Auction):

1. The bidders who have cleared /qualified for the technical evaluation are only allowed to participate in this stage.
2. It is a Techno-Commercial evaluation hence bids will be awarded to vendor with higher scores.
3. Bidders have to quote for commercial bids also.
4. Electronic Reverse Auction may be carried out only if bids received are not competitive. The details of the financial bidding phase will be announced to the short-listed bidder(s) at a later date. Bidders should quote a single figure which includes all the cost of the project only in INR. The successful bidder will submit the item wise bifurcation of a single figure quoted within 24 hours.
5. Base price for Reverse auction will be determined from the price bids submitted in SRM portal. IIT Bombay reserve the right to further negotiate with the L1 bidder

IV. **Bid Evaluation**

1. The Evaluation process to identify the successful bidder is based on the combined techno-commercial evaluation. The bids received from the bidders will be evaluated by the Technical Committee constituted by the Institute. The decision of the technical committee is final and binding on all bidders.
2. The technical bids are evaluated first. The mandatory conditions mentioned elsewhere must be adhered to and failure of the same will result in disqualification of the bid.
3. Each bidder has to obtain the minimum score under each category of the Technical Evaluation to qualify for opening of the commercial bid.

The scores obtained in 1 to 4 above will be added to arrive at the techno-commercial score of the bidder. The bidder with the highest techno-commercial score will be selected as the L1 bidder.

<i>Sl No.</i>	<i>Description</i>	<i>Max Score</i>	<i>Min Score</i>
Technical Bid Evaluation			
1.	OEM Profile (including reference systems as mentioned in the section “Bidders’ Eligibility Criteria” of this document)	60	20
2.	Solution superiority and technical presentation	95	10
3.	Benchmark Performance	195	100
Commercial bid Evaluation			
4.	Price Evaluation	150	0

Technical Bid Evaluation:

Solution superiority: -

No.	Particulars	Bid	Presentation	Total
1	Total core counts from all CPU only nodes	(proposed core count/the most proposed) x Max mark		15
2	Total memory from all CPU only nodes	(Your proposed quantity/the most proposed) x Max mark		10
3	Difference between RPeak and RMax (in percentage); lesser is better (on CPU only)	(The least deviation proposed/deviation of your proposed count) x Max mark		10
4	Network latency between nodes (worst case and average in respect to RTT and TTFB); lesser is better; OEM has to reproduce the result after installation using standard tools	(The least latency proposed/latency of your proposed count) x Max marks		10
5	Number of accelerator per board	(Your proposed number/the most proposed) x Max mark		10
6	Communication among the accelerator	Dedicated arrangements will score maximum marks. All other like PCI based communication will be awarded 5		10
7	Disk rebuilding time in case of a disk failure	(The least time proposed/time proposed in your solution) x Max marks		10
8	Compactness of the solution	(The least footprint proposed/footprint of the proposed solution) x Max mark		10
9	Energy efficiency	(The least consumption proposed/power consumption by proposed solution) x Max mark		10

Benchmark performance: This score will depend on benchmarking runs of prescribed programs on specific inputs as follows.

Benchmark Programs

Benchmark Programs

Vendors are required to run different benchmark applications (links provided below) in different test case scenarios. The details of input files and test cases are mentioned below:

Application	Version	Where to download	Input Deck	Where to Run
CODE I		https://www.me.iitb.ac.in/~sgopalak/benchmarkOpenMP/laplace_dynamic.c		CPU only
GROMACS	Version 2022.2 released on June 16th, 2022	https://ftp.gromacs.org/gromacs/gromacs-2022.2.tar.gz	https://www.mpinat.mpg.de/benchmark/PEP	CPU and Accelerator
WRF	Version 4.3.3 (bug fix release)	https://github.com/wrf-model/WRF/releases/tag/v4.3.3	Standard "conus2.5km" benchmark (https://www2.mmm.ucar.edu/WG2bench/conus_2.5_v3/), with extended length of forecast and 1 hourly output frequency. This is a single domain benchmark with 129 million grid points and grid resolution of 2.5km. This should scale up to 32 nodes or more	CPU only
OpenFOAM	Version 2206	https://www.openfoam.com/news/main-news/openfoam	1. pisoFoam – motorbike case with LES. tutorials/incompressible/pisoFoam/LES/	CPU Only

		m-v2206	2. interPhaseDyMFoam propeller /tutorials/multiphase/interPhaseChangeDyMFoam/propeller	
PETSc		https://www.mcs.anl.gov/petsc/	MPIranks = 10000	CPU and Accelerator

Marks obtained = (The least runtime/projected runtime) x Max score

Max score is 5 for each test cases

Open-MP scalability test

Applications	Runtime (t)			Maximum marks
	Single Core	Single socket	Single board	
Code I				15
Gromacs				15

MPI scalability test (on CPU only nodes): On 1 core per socket

Applications		Runtime (t)					Maximum marks
		1 Node	4 Nodes	8 Nodes	16 Nodes	32 Nodes	
GROMACS							25
WRF							25
Open Foam	pisoFoam						25
	interPhaseDyMFoam						25
PETSc							25

Accelerator scalability test

Applications	Runtime (t)				Maximum marks
	1 Cards	4 Cards	8 Cards	16 Cards	
GROMACS					20
PETSc					20

The total vendor score is a measure of the amount of work that IIT Bombay can expect from the proposed system for each benchmark.

Acceptance criteria:

1. All applications have to be run on machines manufactured and owned by the OEM.
2. All runs should be performed on the processors and system configuration being proposed. If this is not possible, then extrapolations to the proposed configuration are permitted. However, the vendor MUST provide results for all benchmarks generated on hardware that is as close as possible to one that is proposed.
3. Runs can be made with turbo boost and SMT/HT enabled but boost and SMT/HT should be used consistently across all the benchmarks. The vendor must state clearly the settings used.
4. Complete configuration of the systems used for running these applications should be provided along with the output, log files, build and compiler options used, runtime and any other relevant information. Lack of clarity will impact the final technical score.
5. The bidder must provide an explanation of the methodology used to generate extrapolated performance (if any) and numbers. IIT Bombay reserves the right to reject invalid extrapolations. Lack of clarity will impact the final technical score.
6. At the time of delivery and system acceptance, the vendor must execute the applications as per all test cases mentioned in the RFP and produce the result projected in the technical bid. A margin of 5% is accepted.
7. All runs must run to completion successfully with no errors at the time of acceptance test.
8. The bidder must ensure that results of all WRF runs are reproducible (i.e., identical runs given identical output).

OEM Profile: OEM will be judged on the basis of their previous installations in India and abroad (OEM must produce purchase order and acceptance certificate issued by the user - and not by the installation agency or consultant), location, strength and capabilities of the engineering/service team of the OEM, market reputation and the share in the top500.org list. The bidder should submit documents describing their profile, and supporting their claims, as part of the technical bid, else the bid will be rejected.

1	Number of years the OEM has been engaged in HPC rack-based cluster business (supply, installation, support and maintenance. It will be counted from the date of supplied PO in the technical bid, irrespective of private or public sectors) in India. PO must be attached to prove bidders' claims.	>10 yrs will be awarded 10 points. 5 to 10 yrs will be awarded 7 points 3 to 5 yrs will be awarded 4 points 1 to 3yrs will be awarded 2 point 1yr will be awarded 1 point	10
2	Number of HPC rack-based cluster (≥ 50 TF) installations and supply executed by the bidder/OEM (as on date of submission including public and private sectors) in India. PO must be attached to prove bidders' claims.	>10 will be awarded 10 points. 8-10 will be awarded 8points 5-7 will be awarded 6 points 3-4 will be awarded 4 points 1-2 will be awarded 2 point 0 will be awarded 0 point	10
3	Cumulative turnover of OEM in last three financial years from Hardware supply, maintenance and Support service activities of rack based HPC servers only in India. Year 2018-2019 Year 2019-2020 Year 2020-2021 Supporting documentary evidence must be provided.	>500Cr will be awarded 10 points. 350 to 500 Cr will be awarded 8 points 200 to 350 Cr will be awarded 6 points 50 to 200 Cr will be awarded 4 points < 50 Cr will be awarded 2 point Cr= Crore INR	10

4	<p>The OEM should have adequate documented experience in setting up three one HPC clusters capable of at least two 500+ TeraFLOPs (CPU only RMax RPeak) or three HPC clusters capable of at least 250+ TeraFLOPs (CPU only RMax RPeak) each or four HPC cluster capable of at least 150+ TeraFLOPs (CPU only RPeak) each in the last 5 years (in the last 5 consecutive financial years). Purchase order copies along with acceptance reports issued by the end user to the OEM have to be submitted. Also Reference to publicly available evidence like India super-computing site (http://topSC.in) must be provided for verification.</p>	<p>>=2 HPCs containing peak speed of >=500-TF (for CPU only) will be awarded 10 points. 3 HPCs containing sustained speed of >=250 TF (for CPU only) will be awarded 9 points 4 HPCs containing sustained speed of >=150 TF (for CPU only) will be awarded 8 Points Anything lesser than that will not be considered</p>	10
5	<p>The server OEM or the storage OEM (in case server OEM and storage OEM are not same) must have supplied at least 1 number of PFS storage in the country of 1 PiB or 2 number of PFS storage of 500TB or more capacity as part of HPC solution (standalone storage installation will not be counted) in the last 5 years (in the last 5 consecutive financial years). Purchase order copies with bidder installations reports to be submitted along with contact nos. of the person.</p>	<p>>=1 number of PFS storage in the country of >=1 PiB will be awarded 10 points. >=2 number of PFS storage in the country of >=500 TB will be awarded 8 points Anything lesser than that will not be considered</p>	10
6	<p>HPC OEM presence in Top500.org latest 4 reports (Only distinct entry will be counted)</p>	<p>>=10 entries of OEM will be awarded 10 points. >=7 entries of OEM will be awarded 7 points >= 4 entries of OEM will be awarded 4 points >= 1 entries of OEM will be awarded 2 points</p>	10

Commercial Bid Evaluation:

Price Evaluation: Marks are computed based on the **total cost of ownership** (TCO). This is the sum of the cost of the proposed solution which includes HPC cluster with the minimum number of compute nodes specified earlier, storage, interconnect, software, etc.

The marks of a particular bidder are computed as follows:

$$m_c = 150 - [\text{TCO}(\text{bidder}) - \text{TCO}(\text{lowest})] \times 10$$

Where, TCO is in **crores** and TCO (bidder) is TCO of the bidder and TCO (lowest) is the lowest TCO of all the bidders.

I. Bidder's Eligibility Criteria:

1. The bidder has to be an Original Equipment Manufacturer (OEM) of the server hardware.
2. OEM (of server, storage and network peripherals; if they are not the same) must have an active presence in the Indian and global market and experience in supply, installation, support and maintenance of dense HPC clusters with parallel file systems and 100 percent not blocking network arrangements. PO must be attached as a proof of evidence. Years of experience will carry weightage in the evaluation process.
3. OEM will not be considered as Make in India, unless OEM can confirm and prove that the proposed solution is already being manufactured at the time of bidding at the Indian manufacturing unit of the OEM.
4. The bidding OEM brand should have at least 10 entries (in last 4 reports) in Top500 and also must have adequate documented experience during the last 10 years in setting up two HPC clusters worldwide capable of at least 5 PeatFLOPS (CPU only RPeak) each (Only distinct entry will be counted and the number of entries will have an impact in the evaluation process)
5. The OEM should have adequate documented experience in setting up two HPC clusters capable of at least 500+ TeraFLOPs (CPU only RPeak) or three HPC clusters capable of at least 250+ Teraflops (CPU only RPeak) each or four HPC cluster capable of at least 150+ TeraFLOPs (CPU only RPeak) each in India in last 10 years. If desired, the bidder should arrange the visit to such installations by IIT Bombay representatives. Reference to publicly available evidence like india supercomputing site (<http://topSC.in>) along with OEM purchase order and acceptance certificate issued by the user - and not by the installation agency or consultant. Number of entries and size will carry weightage in the evaluation process
6. The OEM (if the storage is outsourced from different OEM, then it applies to the storage OEM as well) should have one installation of parallel file systems of a capacity of at least 1+ PetaBytes (PiB) or two parallel file system of a capacity of at least 500 TB+ (usable) in India in last 10 years with similar product that is being proposed for IIT Bombay. OEM has to attach purchase order and acceptance certificate issued by the user - and not by the installation agency or consultant. Number of entries and size will carry weightage in the evaluation process
7. The bidding OEM should be in a position to deliver a complete end-to-end solution. If any part of the solution is outsourced from another OEM / developer, the bidding OEM has to take responsibility for the complete solution in terms of post sale support and maintenance. An authorization letter from that OEM of the outsourced product has to be submitted with the bid by bidding OEM. However, the bidding OEM has to show those items in the proposed BOQ with the bidding OEM's HSN number. Proposal has to clearly reflect the competence and capability of the bidding OEM in meeting this requirement.
8. The bidding OEM should submit documents showing a minimum annual sales turnover of at least Rs 100 Crores during the last three consecutive years.
9. The bidding OEM should produce a solvency certificate (not older than a year) for Rs 10 Crores issued by a scheduled/nationalized bank with which the bidder holds a current account.
10. An undertaking (self certificate) is to be submitted by the bidding OEM that the organization has not been blacklisted by any Central/State Government Department/Organization and educational institutes in the last three years.

11. The bidding OEM must have local support in Mumbai in terms of service and spare parts. The OEM must already have an operational spares center / warehouse / support office in Mumbai for support services. (setting up a new one for this tender will not be counted)
12. All components Servers/Storage/Network should be shipped in original packing with OEM part numbers mentioned. Bill of material with OEM part nos. should be submitted with technical bid.
13. Canvassing in any form would disqualify the OEM from further participation.

Annexure - 1

Site compatibility proforma

Item	Description of Existing Facility	Whether Compatible with the Proposed Solution?		Remarks
		Yes	No	
False Floor (layout, point load, and height)				
Power Cable (from DB to Socket)				
Power Socket (near the racks)				
Chilled water - quantity, flow rate, water temperature				
Tapping point dimension				
Others (describe in detail)				

Annexure - 2

Performance Statement proforma (for a period of last three years)

Name of the firm

Order Placed by (full name and address of the purchaser)	Order number and date	Description and quality of the ordered equipment	Value of order	Date of completion of delivery as per contract	Date of actual completion of delivery	Reason of late delivery if any	Has the instrument been installed properly? (submit a certificate from the purchaser)	Contact person along with contact details

Signature and seal of the manufacturer:

Date:

Place:

Annexure - 3

Certificate Of Completed Work From Past Customers

(Furnish this information for each individual work from the CUSTOMERS referred in previous Form for whom the work was executed)

1. Name of work / Project and Location
2. Agreement/Purchase Order Number
3. Estimated Cost
4. Tendered Cost
5. Date of Start
6. Date of Completion
 - a. Stipulated date of completion
 - b. Actual date of completion
7. Amount of compensation levied for delayed completion if any.
8. Performance on HPL Benchmark using CPU cores of HPC system (in TFlop/s)
9. Performance report:
 - a) Quality of Work: Excellent/ Very good/ Good/ Fair
 - b) Resourcefulness: Excellent/ Very good/ Good/ Fair
 - c) Responsiveness: Excellent/ Very good/ Good/ Fair
 - d) Accessibility to management when needed: Excellent/ Very good/ Good/ Fair
10. Name of Institute/ Chief Project Manager or Equivalent
11. Contact Details
12. Would you award work again to this supplier Yes/ No

Date:

Place:

Signature (with Seal)