Pre-Bid Meeting Minutes

Tender No.: AMPRI/PUR/EQP/98/ (2025-26)

Subject: Supply, Installation, and Commissioning of Multi-GPU Server with Rack and Network-Enabled Nodes with Supporting Software

Date: 28th January 2026

Time: 2:30 PM onwards

A Pre-Bid Meeting was convened on 28th January 2026 at 2:30 PM onwards at CSIR-AMPRI, Bhopal, to discuss and clarify the technical specifications and bid conditions for Tender No. AMPRI/PUR/EQP/98/(2025-26) pertaining to the procurement of a Multi-GPU Server with Rack and Network-Enabled Nodes with Supporting Software.

The pre-bid meeting was conducted in hybrid mode to facilitate participation from interested bidders both physically and through an online platform. A total of six (6) firms expressed interest in participating in the pre-bid meeting. Out of these, four (4) firms attended the meeting through online mode via a shared video conference link (attendance attached).

List of Firms Attended the Pre-Bid Meeting

Sl. No.	Firm Name	Mode of Participation
1	M/s Daffodil Technologies India Pvt. Ltd.	Online
2	M/s NetwebIndia Pvt. Ltd.	Online
3	M/s CCS Computers Pvt. Ltd.	Online and Physical
4	M/s Meganet Technologies Global Limited	Online

Firms Expressed Interest but Did Not Attend:

The following firms were provided with the online meeting link, but did not join the pre-bid meeting:

- M/s Supremesols
- M/s Inflow Technologies Pvt. Ltd.

The queries raised by prospective bidders were received via email before and during the pre-bid meeting. The detailed responses, recommendations, and justifications are attached as ANNEXURE-A

ANNEXURE-A

1. M/S Daffodil Technologies India Pvt. Ltd.

S. No.	Page No	Sec. No/SR No.	Asked Specification	Queries raised	Response and recommendations
1	46	1.5) Disks Bays	At least 12 LFF or more SATA HDD/SSD Bays.	In 4U rack GPU server, it is a normal and practice to use the HDD's/SSDs of 2.5", due to the heat generation issue in 3.5" HDDs, hence we recommend to make this clause amend to "At least 12 SFF/LFF or more SATA HDD/SSD Bays."	**Accepted with enhancement.** Clause amended to "At least 20 x 2.5" SFF SSD/NVMe drive bays or more" to ensure optimal thermal design, cooling efficiency, and high-performance storage for compute-intensive AI/ML workloads at CSIR-AMPRI. This supports large dataset processing, model training, and real-time inference requirements.
2	46	1.8) I/O slots	10xPCIe Gen 5.0 or more, 8x16 PCIe5.0 FHFL (Double Width) or more, 1x8 PCIe5.0 FHHL (Single Width) or more, 1 x8 PCIe5.0 OCP NIC 3.0 or more	In 4U rack server max. number slots present would be "10xPCIe Gen 5.0 or more, 1x8 PCIe5.0 FHHL (Single Width) or more, 1 x8 PCIe5.0 OCP NIC 3.0 or more	**Accepted with modification.** Clause revised to: "10x PCIe Gen 5.0 or more (of which 8x16 PCIe5.0 FHFL (Double Width) or more, 1x8 PCIe5.0 FHHL (Single Width) or more, 1x8 PCIe5.0 OCP NIC 3.0 or more)" to reflect standard 4U GPU server PCIe slot configuration while maintaining GPU acceleration and networking requirements.
3	46	1.10) Other ports	1 x Display port DB15 (Rear)	Normally The GPU Servers come default with VGA port for display with DB15, so please let us know whether you require VGA port or DP port using any GPU card (through an adapter)?	**Clarified and retained.** Specification refers to 1 x VGA port with DB15 connector (Rear), which is standard for server BMC/iDRAC remote management and initial setup. This is required for proper server accessing during installation, troubleshooting, and headless operation. No amendment needed; retained as original for standard server compliance.
4	47	3.4) Accelerator	Nvidia Quadro RTX 4000,20GB	There is no such GPU card is available. Either it is RTX A4000 GPU which were available with 16GB but not its EOL now. And gets replaced by new RTX Pro 4000 Blackwell GPU card 24GB	**Accepted.** Clause amended to "NVIDIA RTX PRO 4000 Blackwell 24GB or better" to reflect currently available professional workstation GPUs suitable for AI/ML workloads. This ensures vendor compliance with the latest NVIDIA GB200/Blackwell architecture while meeting memory and performance requirements.

5	48	5.1.6)	The Cluster Management Suite must have Web based tool for the administration of HPC, including real-time monitoring, a historical repository based on jobs, users, application etc.	We can offer Open-HPC, which is an open-source, command-line–based cluster management software suitable for provisioning, managing, and monitoring cluster nodes.	**Accepted with conditions.** Open-HPC is approved as the cluster management solution, provided the OEM supplies comprehensive support, including installation, configuration, training, and maintenance throughout the warranty period of the entire system. Web-based monitoring must be enabled via integration with compatible tools (e.g., Ganglia, Prometheus + Grafana).
6	48	5.1)	Workload Manager	The features mentioned are primarily available in GUI-based and paid workload managers. However, we can propose an open-source SLURM workload manager, which fulfils the required workload management functionalities through a command-line interface.	**No change required.** Bidders must provide Workload Manager software exactly matching all functionalities specified in Section 5.1 of the tender document. Licensing & Support Conditions: - Licensed software: License must be in the name of "Director, CSIR-AMPRI, Bhopal" - Open-source OR paid software: OEM must provide complete support (installation, configuration, training, maintenance) throughout the warranty period of the entire system Original specification stands. Compliance will be verified during the technical evaluation
7	48			Please let us know if a GUI-based solution is mandatory or if an open-source CLI-based solution is acceptable for these requirements.	**No change required.** Bidders must provide Workload Manager software exactly matching all functionalities specified in Section 5.1 of the tender document. Licensing & Support Conditions: - Licensed software: License must be in the name of "Director, CSIR-AMPRI, Bhopal" - Open-source OR paid software: OEM must provide complete support (installation, configuration, training, maintenance) throughout the warranty period of the entire system Original specification stands. Compliance will be verified during the technical evaluation

2. M/s NetwebIndia Pvt. Ltd.

High-end GPU Server				
Sr. No.		***Asked Specification***	***Queries raised***	**Response and recommendations**

1.5	Drive Bays	At least 12 LFF or more SATA HDD/SSD Bays	GPU-dense systems typically do not support 3.5" hard drives, as they restrict internal airflow and impact thermal efficiency. Hence, 2.5" drive bays are recommended to ensure optimal cooling. For compute-intensive AI/ML workloads, the system should also provide sufficient atleast 20 or more 2.5" SSD/NVMe drive bays to enable storage scalability, high I/O throughput, and faster data access for improved overall performance.	**Accepted with enhancement.** Clause amended to "At least 20 x 2.5" SFF SSD/NVMe drive bays or more" to ensure optimal thermal design, cooling efficiency, and high-performance storage for compute-intensive AI/ML workloads at CSIR-AMPRI. This supports large dataset processing, model training, and real-time inference requirements.
1.7.	Accelerator	Nvidia H200 NVL (141GB) with Nvlink	NVlink Bridge. https://developer.nvidia.com/blog/deploying-nvidia-h200-nvl-at-scale-with-new-enterprise-reference-architecture/	**Accepted as an upgrade option.** Bidders may offer NVIDIA H200 NVL 141GB (PCIe/SXM) with NVLink bridge support (up to 900GB/s GPU-to-GPU) or equivalent/better, provided it meets or exceeds original performance requirements. H200 NVL delivers 1.5x memory (141GB HBM3e), 4.8TB/s bandwidth, and full NVLink scalability for multi-GPU AI training/inference at CSIR-AMPRI. Must include NVLink bridges, cabling, and OEM validation. Will be evaluated on the technical bid.
		AI/ML Frameworks		**Not required in technical specifications.** The mentioned software libraries, compilers, AI frameworks, inference tools, container platforms, and monitoring utilities are NOT mandatory requirements for technical qualification. Bidders may provide any or all suggested functionalities at no extra
		Library and Compilers Preloaded	GCC, Compilation Environment Setup	
		AI Frameworks Preloaded	1.GPU Optimised Tensor Flow, Pytorch, Theano, Caffe, Text2speech, Mxnet, CuDNN, Containerizing Different Framework, Keras	
			2. Accelerator Optimised Tensor Flow, Pytorch, Theano, Caffe, Text2speech, Mxnet, CuDNN	
		Inference Deployment & Performance Tools	TensorRT	

		Container Platform	System should be preloaded with Precompiled frameworks (CPU & GPU optimized MxNet, CuDNN, Caffe and Pytorch) to be supplied with the system. Data sheet of utility must be submitted with bid.	cost as value addition [illegible] installation/commissioning if th choose. Core hardwar specifications and essential cluster management/workload manager remain the binding requirements.
		Monitoring Tool	The system shall be supplied with a job scheduling utility alongwith resource monitoring & reporting tool.	
1.8.	**I/O slots**	10 xPCIe Gen 5.0 or more, 8 x 16 PCIe5.0 FHFL (Double Width) or more, 1x8 PCIe5.0 FHHL (Single Width) or more, 1 x8 PCIe5.0 OCP NIC 3.0 or more	10 xPCIe Gen 5.0 or more (8 x 16 PCIe5.0 FHFL (Double Width) or more, 1x8 PCIe5.0 FHHL (Single Width) or more, 1 x8 PCIe5.0 OCP NIC 3.0 or more)	**Accepted with modification.** Clause revised to: "10x PCIe Gen 5.0 or more (of which 8x16 PCIe5.0 FHFL (Double Width) or more, 1x8 PCIe5.0 FHHL (Single Width) or more, 1x8 PCIe5.0 OCP NIC 3.0 or more)" to reflect standard 4U GPU server PCIe slot configuration while maintaining GPU acceleration and networking requirements.
5.1.6		**Cluster Manager Software Features:**	Cluster Management Software: The proposed solution shall include an OEM-deployed and OEM-supported cluster management software for a minimum period of five years with a web-based GUI, remote access, role-based access control, and GUI-based management of users, roles, and compute nodes (add/modify/delete). The solution shall be compatible with the latest releases of RHEL and Rocky Linux and shall support OS and application auto-provisioning, profile-based automated provisioning, GUI-based cluster monitoring, and network-wide access to the HPC management console. All required licenses shall be included.	**Not required in technical specifications.** The mentioned specifications related to cluster management software, compilers, libraries, and scientific applications do not directly affect the basic hardware functionality or core performance parameters of the bid system. These features primarily pertain to software-level management and user convenience rather than the underlying computational performance of the proposed hardware. If the bidder wishes, they may provide these functionalities at no additional cost; however, these shall not form part of the technical
			Libraries, MPI, and Compilers: The scope shall include installation and configuration of OpenMP, MPI, C/C++ and FORTRAN compilers, GSL libraries, and processor-optimized BLAS, LAPACK, FFTW, and Fortran compilers. OpenMP and MPI networks shall be implemented, and MPI communication shall be configured over both RoCE and TCP/IP. All compiler and library licenses, including upgrades during the warranty period, shall be provided by the bidder.	

			Scientific & research Applications: The Solution shall install the latest stable version of Rocky Linux, configure the job scheduler, and install common scientific applications such as LAMMPS, Gromacs, VASP, Turbomole, QuantumATK, Gnuplot, Plot Digitizer, and Ovito. Licenses for these applications shall be provided by the user. Cluster functionality testing shall be performed.	evaluation criteria. Since these features do not directly influence the system's core hardware performance or benchmark capabilities, their inclusion or absence will not impact the overall technical assessment of the bid.
		NVLINK Bridge Details	https://developer.nvidia.com/blog/deploying-nvidia-h200-nvl-at-scale-with-new-enterprise-reference-architecture/	**Accepted as an upgrade option.** Bidders may offer NVIDIA H200 NVL 141GB (PCIe/SXM) with NVLink bridge support (up to 900GB/s GPU-to-GPU) or equivalent/better, provided it meets or exceeds original performance requirements. H200 NVL delivers 1.5x memory (141GB HBM3e), 4.8TB/s bandwidth, and full NVLink scalability for multi-GPU AI training/inference at CSIR-AMPRI. Must include NVLink bridges, cabling, and OEM validation. Will be evaluated on the technical bid.

		Annexure-XIII		
		Delivery Schedule	Request to change 120 days instead of 30 days from the purchase order. Due to global shortages of major components.	**Accepted** Please fill in Chapter 3 of the bid document with proper justification and clarification

3. **CCS Computers Pvt. Ltd**

Master System				
Sr. No.		**Asked Specification**	**Queries raised**	**Response**
1.	Performance Benchmark		1.Specrate2017_fp_base >460 2.Specrate2017_int_base >330 System OEM must have listed SPEC benchmark score as afore mentioned in www.spec.org with the same System model or a	**Not required.** SPECrate2017 is irrelevant for GPU-centric HPC; validation via HPL/GPU benchmarks during technical testing. Optional submission; excluded from evaluation

			model from the similar series, with same CPU configuration as quoted in the bid.	as it does not impact core system performance AI/ML workloads.
2.	OS Certification		Quoted Model must be listed on OS official website of Window, RHEL Linux and certificate copy must be attached with bid.	**Not required.** Compatibility verified via live OS installation/testing during technical evaluation (RHEL/Rocky Linux + CUDA drivers). Windows is irrelevant for HPC. Optional documentation; excluded from evaluation as it does not impact hardware/OS integration performance.
GPU System				
3.	Performance Benchmark		1.Specrate2017_fp_base >670	**Not required.** SPECrate2017 is irrelevant for GPU-centric HPC; validation via HPL/GPU benchmarks during technical testing. Optional submission; excluded from evaluation as it does not impact core system performance for AI/ML workloads.
4.	OEM Listing		2.Specrate2017_int_base >500	**Not required.** SPECrate2017 is irrelevant for GPU-centric HPC; validation via HPL/GPU benchmarks during technical testing. Optional submission; excluded from evaluation as it does not impact core system performance for AI/ML workloads.
5.	OS Certification		System OEM must have listed SPEC benchmark score as afore mentioned in www.spec.org with the same System model or a model from the similar series, with same CPU configuration as quoted in the bid.	**Not required.** SPEC irrelevant to OS certification (RHEL/Rocky Linux verified via OEM HCL and live testing). Optional; excluded from evaluation as it does not impact OS/hardware integration or performance.
Workstation				
6.	System utility with AI Frameworks Preloaded		System should be preloaded with Precompiled frameworks (CPU & GPU optimized MxNet, CuDNN, Caffe and Pytorch) to be supplied with the system, license must be in named of the Customer organization, Data sheet of utility must be submitted with bid.	**Not required.** Frameworks available open-source; optional bidder value-add at no extra cost. Excluded from technical evaluation as it does not affect hardware performance.
7.	Job Scheduling Utility License		Unified system management/monitoring toolset for configuration, diagnosis and management of the system. Toolset/Manager must be capable of supporting package and image-based provisioning, intuitive web interface for managing and customize the node, And tool set with provisioning, monitoring, and reporting capabilities. With JOB scheduling capabilities on single node for CLI and GUI-based end-user	**Not required.** SLURM/OpenHPC covers all functionalities (CLI/GUI via plugins like Slurm Web/Grafana). Optional at no extra cost; excluded from evaluation as non-impactful to hardware performance.

			applications. S/W utility must be from the H/W OEM and not from a third party vendor. The product datasheet with demonstration details must be shared	
8.	OEM Listing		scheduling capabilities on single node for CLI and GUI based end user applications. S/W utility must be from the H/W OEM and not from a third-party vendor. Product datasheet with demonstration details must be shared	**Not required.** SLURM provides equivalent CLI/GUI scheduling; OEM restriction unnecessary. Optional at no extra cost; excluded from evaluation as it does not impact system performance.
9.	BIS Registration		The Server OEM must have its systems listed with any Blackwell GPU variant in the NVIDIA Qualified Systems Catalogue.	**Not required.** GPU compatibility verified via CUDA drivers and OEM benchmarks. Optional documentation; excluded from evaluation as it does not impact hardware compute performance.
10.	Regulatory Compliance		OEM must have Bureau of Indian Standards certification for Workstation product category and necessary document must be provide.	**Not required.** BIS certification is regulatory documentation, not performance-related. Verified via OEM standard certifications and benchmarks. Optional submission; excluded from evaluation as it does not impact hardware performance.
Eligibility Criteria				
11.			OEM must have its own registered Service and support center in Delhi. Details must be listed on the OEM Website. And OEM must provide documentary evidence issued by Govt department.	**Not required.** On-site support mandated via SLA with defined response times; Delhi-specific requirement unnecessary and discriminatory. Pan-India OEM support networks are sufficient. Compliance verified through SLA execution record.

4. M/s Meganet Technologies Global Limited

S. N	Page No	Sec.No /SR NO	Asked Specification	Queries raised	Response and recommendations
1	46	1.5) Disks Bays	At least 12 LFF or more SATA HDD/SSD Bays.	In 4U rack GPU server, it is a normal and practice to use the HDD's/SSDs of 2.5", due to the heat generation issue in 3.5" HDDs, hence we recommend to make this clause amend to "At least 12 SFF/LFF or more SATA HDD/SSD Bays"	**Accepted.** Clause amended to "At least 12 SFF/LFF or more SATA HDD/SSD Bays" to accommodate standard 4U GPU server thermal design and performance requirements.

2	46	1.8) I/O slots	10xPCIe Gen 5.0 or more, 8x16 PCIe5.0 FHFL (Double Width) or more, 1x8 PCIe5.0 FHHL (Single Width) or more, 1 x8 PCIe5.0 OCP NIC 3.0 or more	In 4U rack server max. number slots present would be "10xPCIe Gen 5.0 or more, 1x8 PCIe5.0 FHHL (Single Width) or more, 1 x8 PCIe5.0 OCP NIC 3.0 or more	**Accepted with modification.** Clause revised to: "10x PCIe Gen 5.0 or more (of which 8x16 PCIe5.0 FHFL (Double Width) or more, 1x8 PCIe5.0 FHHL (Single Width) or more, 1x8 PCIe5.0 OCP NIC 3.0 or more)" to reflect standard 4U GPU server PCIe slot configuration while maintaining GPU acceleration and networking requirements.
3	46	1.10) Other ports	1 x Display port DB15 (Rear)	Normally The GPU Servers come default with VGA port for display with DB15, so please let us whether you require VGA port or DP port using any GPU card (through adapter)?	**Clarified and retained.** Specification refers to 1 x VGA port with DB15 connector (Rear), which is standard for server BMC/iDRAC remote management and initial setup. This is required for proper server accessing during installation, troubleshooting, and headless operation. No amendment needed; retained as original for standard server compliance.
4	47	3.4) Accelerator	Nvidia Quadro RTX A4000,20GB	There is no such GPU card is available. Either it is RTX A4000 GPU which were available with 16GB but not its EOL now. And gets replaced by new RTX Pro 4000 Blackwell GPU card 24GB	**Accepted.** Clause amended to "NVIDIA RTX PRO 4000 Blackwell 24GB or better" to reflect currently available professional workstation GPUs suitable for AI/ML workloads. This ensures vendor compliance with the latest NVIDIA GB200/Blackwell architecture while meeting memory and performance requirements.
5	48	5.1.6)	The Cluster Management Suite must have Web based tool for the administration of HPC, including real-time monitoring, a historical repository based on jobs, users, application etc.	we can offer Open-HPC which is an open-source, command-line–based cluster management software suitable for provisioning, managing, and monitoring cluster nodes.	**Accepted with conditions.** Open-HPC is approved as the cluster management solution, provided the OEM supplies comprehensive support, including installation, configuration, training, and maintenance throughout the warranty period of the entire system. Web-based monitoring must be enabled via integration with compatible tools (e.g., Ganglia, Prometheus + Grafana).
6	48	5.1)	Workload Manager	The features mentioned are primarily available in GUI-based and paid workload managers. However, we can propose an open-source SLURM workload manager, which fulfils the required workload management functionalities through a command-line interface.	**Not required.** Bidders must provide Workload Manager software exactly matching all functionalities specified in Section 5.1 of the tender document. Licensing & Support Conditions: - Licensed software: License must be in the name of "Director, CSIR-AMPRI, Bhopal" - Open-source OR paid software: OEM must provide complete support (installation, configuration, training, maintenance) throughout the warranty period of the entire system

ANNEXURE-A

					Original specification stands. Compliance will be verified during the technical evaluation
7	48			Please let us know if a GUI-based solution is mandatory or if an open-source CLI-based solution is acceptable for these requirements.	**Not Required** Bidders must provide Workload Manager software exactly matching all functionalities specified in Section 5.1 of the tender document. Licensing & Support Conditions: - Licensed software: License must be in the name of "Director, CSIR-AMPRI, Bhopal" - Open-source OR paid software: OEM must provide complete support (installation, configuration, training, maintenance) throughout the warranty period of the entire system Original specification stands. Compliance will be verified during the technical evaluation

Note: All other specifications in the bid document remain unchanged.

HIGH-PERFORMANCE COMPUTING (HPC) CLUSTER FOR CSIR-AMPRI, BHOPAL
TECHNICAL SPECIFICATIONS (Revised)

1.	**High-end GPU Server - Qty 1**	
1.1.	**Form Factor**	Max. 4U (including all units) rack-mounted with sliding rails
1.2.	**Configured CPUs**	Min Two or more x86 Architecture-based Processors, each CPU with 24+ Cores, 3.2GHz Base or higher, with 128MB L3 Cache or more, and optimized for GPU offload
1.3.	**NVMe Support**	At least one or more drives/Slots must have native support for NVMe drives.
1.4.	**Memory configured**	1024GB DDR5 5200 MHz+ Registered ECC RAM installed from day one. Total 24+ DIMMS Slots
1.5.	**Disks Bays**	At least 20 x 2.5" SFF SSD/NVMe drive bays or more
1.6.	**SSD & HDD**	1 x 3840 GB or higher Ent. M.2 NVMe SSD Drives installed from day one.
1.7.	**Accelerator**	2 x Nvidia H200 NVL (141GB) with Nvlink Bridge System must support eight such accelerators from day one, or at least 8 double-width or lesser-width based GPUs with required expansion slots and optimum power supplies
1.8.	**I/O slots**	10x PCIe Gen 5.0 or more (of which 8x16 PCIe5.0 FHFL (Double Width) or more, 1x8 PCIe5.0 FHHL (Single Width) or more, 1x8 PCIe5.0 OCP NIC 3.0 or more)
1.9.	**Interconnect**	2 x 1Gbps Ethernet Ports with through onboard or add-on Controller. 2 x 10G Base T LAN Ports
1.10.	**Other ports**	1 x 1000Base-T Dedicated Server Management Port (Rear), 2 x 1GBase-T Ethernet ports with Intel® i350 (Rear), 2 x USB 3.2 Gen1 Type-A Port (Rear), 4 x USB 3.2 Gen1 Type-A Port (Front), 1 x Display port DB15 VGA (Rear)
1.11.	**Power Supply**	Redundant N+1 or more, 2700W 80+Titanium Redundant Power Supplies.
1.12.	**Installation**	Installation, Testing, Training, and Implementation costs for all above mentioned solution must be included from day one. (Done by OEM certified Engineers)
1.13.	**Management**	IPMI 2.0 with virtual media over LAN and KVM-over-LAN support.
1.14.	**Security**	Secure Platform Module, Hardware-Based Security, Secure Key Generation and Storage, Device Authentication/Remote Attestation, Integration with OS Features, Platform Integrity (Measured Boot).
1.15.	**Warranty**	Minimum 5 years of comprehensive Part, Labor, and Onsite service by OEM certified engineer.
2.	**Master System – 1 Qty**	
2.1.	**Form Factor**	Max. 2U (including all units) rack-mounted with sliding rails
2.2.	**Configured CPUs**	Min Two or more x86 Architecture-based Processors, each CPU with 16+ Cores, 2.6GHz Base or higher, with 64MB L3 Cache or more.
2.3.	**NVMe Support**	At least one or more drives/Slots must have native support for NVMe drives.
2.4.	**Memory configured**	512 GB DDR5 5200MHz Registered ECC RAM installed from day one. Total 16 DIMMS Slots or higher
2.5.	**Disks Bays**	At least 12 LFF or more SATA HDD/SSD Bays.
2.6.	**RAID Controller**	RAID Controller SAS3 12Gbps with 1GB Cache and Supports RAID levels 0, 1,5,6,10
2.7.	**SSD & HDD**	2 x 1920GB SATA SSD NVMe Gen5 SSDs 3 x 20000GB or higher Ent. SATA HDD Drives installed from day one.

HIGH-PERFORMANCE COMPUTING (HPC) CLUSTER FOR CSIR-AMPRI, BHOPAL
TECHNICAL SPECIFICATIONS

2.8.	**I/O slots**	2 x16 PCIe4.0 (CPU0),2 x16 PCIe5.0 + 1 x16 PCIe4.0 (CPU1), 3 MCIO PCIe 5.0 x8
2.9.	**Interconnect**	2 x 1Gbps Ethernet Ports with through onboard or add-on Controller. 2 x 10G Base T LAN Ports
2.10.	**Other ports**	2x USB 3.2 (Rear), 1x VGA, 1x dedicated RJ45 Management, 2 x USB 3.0 Type A Ports 2 x USB 2.0 Type-A Ports, 1 x DB-15 VGA connector, 1 x external UID/switch button.
2.11.	**Power Supply**	800W 80+ Redundant Platinum-rated Power Supplies.
2.12.	**Installation**	Installation, Testing, Training, and Implementation costs for all above mentioned solution must be included from day one. (Should be done by OEM certified Engineers)
2.13.	**Management**	IPMI 2.0 with virtual media over LAN and KVM-over-LAN support.
2.14.	**Security**	Secure Platform Module, Hardware-Based Security, Secure Key Generation and Storage, Device Authentication/Remote Attestation, Integration with OS Features, Platform Integrity (Measured Boot).
2.15.	**Warranty**	Minimum 5 years of comprehensive Part, Labor, and Onsite service by the OEM-certified engineer.
3.	**Network-enabled nodes – 2 Qty**	
3.1.	**Processor(s)**	One or more x86 Architecture-based Processors, CPU with 8+ Cores, 3.6GHz Base or higher, with 60MB Cache or more.
3.2.	**Chipset**	Compatible Chipset or better
3.3.	**Memory**	128GB DDR5 Memory or more, by using 32GB memory modules. Scalability after populating with higher-capacity DIMMs must be up to 512 GB or higher.
3.4.	**Accelerator**	NVIDIA RTX PRO 4000 Blackwell 24GB or better
3.5.	**HDD & SSD**	1 x 8000 GB 3.5" SATA 7200 RPM HDD 1 x 960GB M.2 NVMe-based Solid State Drive
3.6.	**NIC**	2 x 1G LAN ports or higher through AOC or through onboard controller
3.7.	**Exp Slots(s)**	2 x PCIe Gen 5.0 x16
3.8.	**Power Supply**	1200W Platinum Level Certified Power Supplies.
3.9.	**Ports**	1 x RJ45 GbE connector for BMC dedicate management port (RTL8211E), 2 x 10GbE (RJ45) by Intel®X710-AT2, 2 x USB 3.0 Type A Ports (Front), 2 x USB3.2 Gen1 type- A Ports (Rear), 1 x DB-15 VGA connector, 1 x external UID/switch button.
3.10.	**Form Factor**	Rack Mount capabilities with up to 5U/Tower or less per node (including all units)
4.	**Rack, UPS, and its monitoring/support/services**	
4.1.	**Primary Switch**	1 x Primary 16 Ports or more 10G Base T LAN Switch with all the required cables
4.2.	**Rack**	1 x Standard 42U Rack
		PDU- 2x7.2kVA
		5 years warranty
4.3.	**Remote monitoring System**	Must be integrated with the system by the bidder/OEM
5.	**Software stack for Job Scheduling**	

ANNEXURE-A

HIGH-PERFORMANCE COMPUTING (HPC) CLUSTER FOR CSIR-AMPRI, BHOPAL
TECHNICAL SPECIFICATIONS

5.1.	**Including Below Features**	
5.1.1.	Batch and interactive jobs: Must support submission, queuing, prioritization, dispatch, and termination of batch and interactive jobs.	
5.1.2.	**Resource management:** Must allocate CPU cores, memory, GPUs, and nodes with enforceable limits per job.	
5.1.3.	Topology-aware placement: Should consider NUMA, sockets, and network topology for optimal job placement.	
5.1.4.	Preemption and checkpointing (suspend/checkpoint/requeue) allow users to immediately run high-priority work	
5.1.5.	Plugin framework for custom health checking, mitigation, and notification capabilities, including offlining flaky Systems, restarting scheduling cycles, and requeuing jobs	
5.1.6.	The Cluster Management Suite must have Web based tool for the administration of HPC, including real-time monitoring, a historical repository based on jobs, users, application etc.	
5.1.7.	Fully featured version with a support-backed suite, with no limitation on the number of jobs and scalability of Systems	
5.2.	Software Support for both Serial and Parallel Environments	Yes
5.3.	Public Documentation: Comprehensive admin and user documentation provided at the time of installation.	
6.	**Supply, Installation, In-house training, and after installation support**	
6.1.	Before acceptance and installation of any item, the quality, specification, and quantity will be verified by CSIR-AMPRI, Bhopal.	
6.2.	Complete Onsite Installation, Testing, Training, Implementation, on-site service, and support must be performed by the OEM's engineer. The scope of work will include supply, racking, OS and application installation, commissioning, and time–to– time servicing of the entire solution being quoted.	
6.3.	Two-day in-house training by Certified Professionals of OEM at CSIR-AMPRI, including installation of software, benchmarking HPL, monitoring of HPC Cluster for 24 Hours, LINUX commands, HPC Management, etc. Scripts for the benchmarking calculations will be provided on request.	
6.4.	After installation, support must be provided by HPC Linux Experts of the bidder on onsite and remote assistance basis during the complete warranty period of 5 years. The bidder must have an internal HPC expert team available for the same.	
6.5.	Bidders must submit a bid-specific Manufacturer Authorization Certificate Form (MAF) on the OEM's Letterhead	
6.6.	All the licenses, if any, should be of a perpetual nature and will be installed by the OEM/bidder.	
6.7.	Any component required for a fully functional solution should be provided along with it at no extra cost	
6.8.	At the time of installation and commissioning of the system by the vendor/OEM, if it is found that some additional hardware or software items are required to make the system fully functional, but not included in the bidder's original list of deliverables, the bidder shall supply such items at no extra cost.	
6.9.	Data transfer at the required speed is demonstrated with proof at the installation site.	
6.10.	Provide end-to-end latency measurements and bandwidth tests for GPU-to-GPU across nodes	
6.11.	Cluster configuration, if any, has to be done by the OEM/Vendor.	
6.12.	Detailed documentation about hardware installation, configuration of the system with a diagram and flowchart, etc., is to be provided at the time of installation.	

HIGH-PERFORMANCE COMPUTING (HPC) CLUSTER FOR CSIR-AMPRI, BHOPAL
TECHNICAL SPECIFICATIONS

6.13.	No component reaches the end of life before others, and the entire system should be provided as a single solution
6.14.	The SITC must be completed on a turnkey basis, integrating hardware, system software, and application/management software (if any) by a single system integrator (certified engineer from OEM)
6.15.	All goods should be new and unused.
6.16.	Bidder shall maintain all the necessary spare parts for faster response.
6.17.	Appropriate hardware nodes/controllers, disks, network switches, and all required accessories and licenses should be included.
6.18.	Irrespective of usage, all hardware should be covered under warranty as specified
6.19.	Faulty hard disks will not be returned to the OEM/Supplier against warranty replacement.
6.20.	Replacement of any part (including disks, nodes, controllers) shall be done with new and genuine parts from the OEM. Refurbished parts should not be provided.
6.21.	Licensed software: License must be in the name of "Director, CSIR-AMPRI, Bhopal."
6.22.	Open-source OR paid software: OEM must provide complete support (installation, configuration, training, maintenance) throughout the warranty period of the entire system.
7.	**Experience and other conditions**
7.1.	The Server OEM should have executed at least 2 HPC Cluster projects in India similar to this tender in the Government Indian academic and research institutions or other Government education & research organizations in India. Details/Proof of the same must be submitted with the technical bid.
7.2.	Server OEM must not have been banned or blacklisted by any government organization in India during the last 5 years
7.3.	OEM and Bidder must not be put on a holiday period or banned or debarred by any govt organization for >3 months during the last 5 years.
7.4.	The quoted model should be BIS certified. All Certificate copies must be attached with the Bid
7.5.	OEM must be ISO 27001, RoHS or similar, Govt. of India E-Waste Regulation Compliance/ Registered. All Certificate copies must be attached with the Bid
7.6.	Make in India System offered must comply with DIPPs notification no. P-45021/2/2017-PP(BE-II) with revisions till date and compliance with P-45021/102/2019-BE-II-Part (1) (E-50310) Office Memorandum issued by GOI. The OEM Certificate should be submitted with the bid