

CSIR- Fourth Paradigm Institute

Expression of Interest (EOI) For Supply, Installation and Maintenance of High-Performance Computing, Artificial Intelligence and Machine Learning Platform (4PI/PUR/2021/34 dated 29th October 2021)

Consolidated Response for Received Queries till 12th November 2021

Sl. No.	Reference	Details	Query	CSIR-4PI Reply
GPU/Accelerator				
1	Pg 8: 1.1.b.i	The system is expected to have at least 8 GPUs per node with total aggregate GPU memory of at least 256 GB and bandwidth of at least 600 GBps for inter GPU communication	<p>Request consider below change: The system is expected to have at least 4 GPUs per node with total aggregate GPU memory of at least 256 GB and bandwidth of at least 600 GBps for inter GPU communication</p> <p>We request that NvLink based 4GPU systems be also included alongside 8 GPU based systems.</p> <p>The system is expected to have at least 8 GPUs per node with total aggregate GPU memory of at least 256 GB and bandwidth of at least 550 GBps for inter GPU communication along with a system memory of about 1TB</p>	<ul style="list-style-type: none"> The system is expected to have at least 4 GPUs per node with per GPU memory of 32 GB or more and total aggregate bandwidths of at least 550 GBps for inter GPU communication. However, in the case of solutions with 4 GPU/node, 8 GPU performance (i.e., across two nodes) of the requested benchmark still needs to be provided in case the bidder is proposing 4 GPU systems.

Sl.No.	Reference	Details	Query	CSIR-4PI Reply
2	Pg. 8: 1.1.b	About one third (1/3)rd of computing power (i.e. 0.75 PetaFlops) should be from accelerator-based nodes suitable for AI workloads.	EOI referred above says the system shall be 2/3rd CPU & 1/3rd GPU based. We request you to consider a portion of it to be IPU based. IPU based system can be integrated seamlessly with other components (CPU compute, GPU compute, storage) of the solution. See the attached reference architecture which is an example.	A bidder can propose IPU as an additional solution in response to this EoI. However, CSIR-4PI will evaluate and considered this proposal if there is merit in the solution subject to demonstrating the benchmark applications on IPU and satisfies the project objective.
3			For accelerator nodes, please define the peak Petaflops performance in terms of FP32/FP64 etc. as the performance varies based on this factor and directly impacts the GPU/Node quantity	About 0.75 PetaFlops of peak performance on FP64 is expected from the accelerator based nodes.
Benchmark suite and benchmark system				
4	Pg. 8: 3	General queries regarding benchmark collection and methods		The benchmark files are only available through authenticated "sftp" site. Hence, each vendor needs to collect them separately after signing the NDA. No other mechanism such as "http" and "https" can be used to download.

Sl.No.	Reference	Details	Query	CSIR-4PI Reply
5		General remark on benchmark		While installing the benchmark the vendor should ensure that they install the desired version of the benchmark (provided in the benchmark document) from their respective sites. Otherwise, there are possibility of errors related to configurations.
6				It is reiterated that successful completion of benchmark and providing the output is one of the essential qualifications to be selected for the next phase.
7	Pg. 29: 5.a.(iii)	Modifications to demonstrable benchmark due to change in the minimum sizing of the GPU node		HPL performance in TF of a GPU node with the proposed configuration.
8	Pg. 29: 5.a.(iv)			MLPerf HPC Training v0.7 (CosmoFlow & DeepCAM) performance of a GPU node with the proposed configuration.

Sl.No.	Reference	Details	Query	CSIR-4PI Reply
9	Pg. 29: 4	(i) SU2 benchmark on CPU only nodes <ul style="list-style-type: none"> • 8 nodes • 16 nodes • 32 nodes • 32 nodes (16 nodes each connected to separate switches) (ii) GROMACS on CPU only nodes <ul style="list-style-type: none"> • 8 nodes • 16 nodes • 32 nodes • 32 nodes (16 nodes each connected to separate switches) 	It is requested to limit the node count to 16 nodes as the bidder may not have a 32-node system with latest CPUs to perform the benchmark. It is also requested to allow extrapolation from earlier generation CPU/Systems	No change to the size of the system. Extrapolation is not a choice of preference. However, extrapolation may be provided from a system of one generation previous to the proposed CPU system only for the purpose of this Eol.
10	Pg. 28: 3	Benchmark System and Ownership Information <ul style="list-style-type: none"> • Configured with a single switch of 100 Gbps • Configured with two switches of 100 Gbps each connecting half of the nodes 	Our benchmark systems are already configured in a specific topology and are used for many worldwide requirements. For 32 node benchmarks, it may not be possible to reconfigure in a way specified in the Eol.	Benchmark System <ul style="list-style-type: none"> • Configured with a single switch of 100 Gbps. • Configured with two switches of 100 Gbps each connecting half of the nodes (if possible). (It may be noted that for benchmark system, which is not reconfigurable, this can also be achieved by manually selecting nodes that are not connected to the same switch)

Sl.No.	Reference	Details	Query	CSIR-4PI Reply
Memory				
11	Pg. 30: 5.b.(iii)	Type of memory (minimum DDR4 3200 MHz ECC RAM)	Instead of minimum DDR4 3200 MHz ECC RAM. Request you to consider below change: Minimum 2.5GBps per core memory bandwidth as measured by STREAM-TRIAD (Collective CPU to memory bandwidth per node / no. of cores in the node)	No change
Storage				
12	Pg. 8: 1.2	A high-performance parallel file system with a raw capacity of about 5 PiB for use as the scratch file system for both GPU and CPU based nodes	<p>Please define the performance throughput (Read/Write GB/s) required in order to correctly size the PFS storage.</p> <p>What is the minimum expected Read/Write throughput?</p> <p>What is the performance requirement expected on the 5PB capacity?</p> <p>We request to clarify what sustained Read/Write performance bandwidth in (GB/s) is required /expected.</p>	The sustained Read/Write performance should be more than 60GB/s.

Sl.No.	Reference	Details	Query	CSIR-4PI Reply
13			Raw capacity is measured in PB/TB/GB/MB/KB (and not PiB/TiB/GiB/MiB/KiB). We request CSIR 4PI to please confirm whether storage requirement is 5PB raw (5,000,000,000,000,000 Bytes) or 5 PiB raw (5,629,499,534,213,120 Bytes)?	The storage requirement is 5 PiB of raw capacity.
14	Pg. 29: 5.a	Demonstrable IoR, HPL and MLPerf Performance	IOR/IoZone performance be limited to a subset of the overall capacity and performance.	Uniform demonstrable IoR performance across the entire storage is essential.
Compliance to Local content				
15	Pg. 32: 16	Local Content of the component & Class of the bidder as per Government of India guidelines.	<p>Typically to determine the class of the bidder, % of local content is calculated as below:</p> <p>Local content % = Total price of Local content / Total price of the solution and not for individual line items.</p> <p>How will CSIR4PI calculate MII (local content)? Please confirm.</p>	CSIR-4PI would strictly follow the public procurements (preference to make-in-India) OMs and guidelines issued in this regard time to time by Government of India, Ministry of Commerce and Industry and the gazette notification dated September 7 2020 by Ministry of Electronics and Information technology (MeitY) notifying electronics product Make-in-India rules. (https://dpiit.gov.in/public-procurements)

Sl.No.	Reference	Details	Query	CSIR-4PI Reply
16				The vendor should provide complete details of the local content value addition under make-in-India including the location(s) at which the local value addition is made. Should also provide certificates from the statutory auditor or cost auditor of the OEM company.
Additional Scope				
17	Pg. 8: 2	Buyback the existing system	Will the buyback value be considered in evaluating L1?	Yes
18	Annexure -A (5.1)	Maintain (warranty and AMC) and provide facility management services for the complete HPC facility	Is the Data Centre currently being supported by any System Integrator OR is the support back-aligned with the OEM?	Yes, The Data Center is currently supported by M/s HCL who has back-to-back arrangement with respective OEMs.
19	Pg. 8: 2	(computing infrastructure + Data center) for 7 years	Additional information on Data Center Components	The vendors are requested/suggested to visit the data center for more detailed information with prior permission.

Sl.No.	Reference	Details	Query	CSIR-4PI Reply
End-of-life				
20	Pg. 31: 12	Undertaking that the quoted product should not be End-of-Life (EoL) for next 7 years.	<ul style="list-style-type: none"> It is possible to provide 3 yrs. warranty and 4 yrs. AMC. It is not possible to provide undertaking on quoted product will not be EoL for the next 7 yrs. 	<ul style="list-style-type: none"> Yes, it is possible to provide 3 years warranty and 4 years AMC. However, the total cost for 7 years will be used for calculating the L1. Provide an undertaking from OEM that the products will be supported for next 7 years
Eol-RFQ restrictions				
21	Pg. 12: 10	<p>a. The OEM (Compute node) can participate through only one SI and vice-versa.</p> <p>b. The OEM-SI pair participated in Eol should be same at the RFQ stage as well.</p>	In case, OEM of major technology components needed to fulfill the solution, wish to respond to the EOI, can we also share the technology details to another System Integrator who wish to submit their response to the EOI	<ul style="list-style-type: none"> No, If OEM is participating directly then SI cannot participate for same major OEM [here compute node OEM] However, at the RFQ stage the OEM may select an appropriate SI (only one, if required) with proper authorization.

Sl.No.	Reference	Details	Query	• CSIR-4PI Reply
22		The proposed solution (in terms of processor, GPU, storage etc.) in the EoI should be the same or next generation while participating in the RFQ. The bidder cannot quote a completely different solution than the one proposed in the EoI.	The final solution and proposal depend largely on the benchmark outcomes that further depends on the various benchmark system availability to derive the best performing system components such as CPU make, Storage throughput etc. Hence, it is requested to allow multiple solutions to be submitted during EOI response for OEMs to have a certain degree of flexibility in positioning the most optimal solution.	<ul style="list-style-type: none"> • The major OEM can propose multiple solutions, if essential, to provide better insight on the available options at the EoI stage. • The bidder needs to fill separate response sheet (Format 9) for each proposed solution. • However, at the RFQ stage one OEM can propose only one solution that the bidder feel most suitable for the requirement and through one SI (if required) only.
Transition plan				
23		<ul style="list-style-type: none"> • The vendor should propose a transition plan for installation of the new system and dismantling of the old system, the time lines etc. in a separate sheet. • Note: CSIR-4PI prefers to install the storage first and move the data from existing DDN storage before dismantling the computing infrastructure 	We would like to seek clarification on whether hard links are activated on the existing DDN storage	No

Sl.No.	Reference	Details	Query	CSIR-4PI Reply
Eligibility to participate				
24	Pg. 11: 6	Bidder eligibility for participating in Eol	Whether consortium is allowed or not? If Yes, how many partners are allowed for forming the consortium?	Consortium is allowed provided consortium is formed and registered prior to bid submission date and for number of partners, refer to the rules and regulations regarding formation of consortium by Govt. of India
25			Do eligibility of consortium partners be considered for meeting the different points as mentioned in eligibility criteria?	Eligibility of consortium partners may be considered for meeting the eligibility criteria
26		Bidder (compute node OEM/SI) should have carried out at least 3 installation.	Need more clarity	<ul style="list-style-type: none"> If a SI is bidding, the SI/Compute node OEM should have carried out minimum 3 HPC installations of desired capacity as mentioned in the Eol document. In such cases where, only the OEM satisfies the eligibility criteria, the OEM needs to provide an undertaking stating that it will take full ownership such as installation, configuration, benchmark demonstration, 7 years 24x7 operation etc. If an OEM is bidding directly, the OEM should have carried out minimum 3 HPC installations of desired capacity mentioned in the Eol document (apart from other eligibility criteria mentioned in the Eol).

Sl.No.	Reference	Details	Query	CSIR-4PI Reply
Additional Information				
27	Pg. 29	Errata	Repetition of serial numbers	<p>Henceforth, a new serial number 5.a in EoI document will refer to Demonstrable IoR, HPL and MLPerf Performance.</p> <p>Likewise, a new serial number 5.b is added to the EoI document to refer CPU based Nodes.</p>
28		General		Bidder should be prepared to make technical presentation of their solution if needed/asked by CSIR-4PI.