Indian Institute of Science (IISc)
Bangalore - 560012

Supercomputer Education and Research Centre (SERC)
IISc

Corrigendum and Response to Queries with reference to
Tender No.: IISc/Purchase/SERC/2020/02 dated February 12, 2020
CPPP Tender ID: 2020_IISC_543758_1

for

Supply and Installation of a multi-TeraFlops High Performance and Heterogeneous Cluster in Supercomputer Education and Research Centre, Indian Institute of Science, Bangalore

Supercomputer Education and Research Centre (SERC)
Indian Institute of Science (IISc)
Bangalore – 560012, India
Email: tender.serc@iisc.ac.in

CPPP Website for e-Tender Submission
https://eprocure.gov.in/eprocure/ap

Contents

1. Corrigendum to the Tender

SNo	Reference	Corrigendum
1	Page 4 1.a. CPU-only-Regular-Nodes SNo 5 NIC for MPI Communications	Modified to: 100 Gbps Infiniband controller Should support hardware packet offload features.
2	Page 4 1.a. CPU-only-Regular-Nodes SNo 9 Form factor	1U/2U Rack mountable Chassis Modified to: Single node chassis of 1U/2U form factor or multiple node chassis with equivalent form factor. Chassis must be rack mountable.
3	Page 5 Master node Form factor	The form factor of the master node can be 3U/4U rack mountable chassis. The rest of the specifications remain the same as in the original tender.
4	Page 5 Master node Storage	Should have extra storage, in addition to the 500 TB, of adequate capacity and capability for storing OS images and network booting of all the nodes in the cluster. This storage should have RAID 6 configuration for reliability. Modified to: Should have extra storage, in addition to the 500 TB, of adequate capacity and capability for storing OS images and network booting of all the nodes in the cluster plus 500 GB max for storing site-specific application binaries. This storage should have RAID 6 configuration for reliability.
5	Page 5 Master node Cross-compilation capability	The master node should provide cross-compilation capabilities for creating executables for the CPU/GPU nodes. If the vendor cannot provide such cross-compilation capability for the GPU V100 nodes, the vendor should provide a storage of 1 TB Enterprise SATA 3.5" @ 7200 RPM in each of the V100 nodes.

6	Page 5 Master node Make and model	Addition of a clause: Note that while the other specifications, **particularly the CPU architecture**, should be the same as that of the other CPU-only-Regular-Nodes, the make and model of the master node can be different from those of the other CPU-only-Regular-Nodes.
7	Page 6 1.d V100 Nodes Form factor	The form factor of the V100 nodes can be 3U/4U rack mountable chassis. The rest of the specifications remain the same as in the original tender.
8	Page 6 1.e SSD Nodes SNo 15	NVMe M2 SSDs of at least 2 TB Modified to: NVMe SSDs of at least 2 TB
9	Page 6 2. Peripherals SNo 1 Interconnect switches and cables	1. Primary interconnect for MPI communications Infiniband switch with appropriate number of ports and cables to connect all the nodes mentioned above. Modified to: 1. Primary interconnect for MPI communications 100 Gbps Infiniband switch with appropriate number of ports and cables to connect all the nodes mentioned above.
10	Page 7 Section 5 Bidder's Eligibility Criteria	Modified to: The bidder should have implemented systems with peak performance of greater than 75 TFlops in at least three sites in the past 5 years. The bidder must provide three references where they have carried out supply, installation and support of systems with peak performance of greater than 75 TFlops in the past 5 years. IISc shall independently obtain inputs from the provided referees before arriving at a final decision. Modified to: The bidder should have implemented systems with peak performance of greater than 75 TFlops from only CPU nodes in at least three sites in the past 5 years. The bidder must provide three references where they have carried out supply, installation and support of systems with peak performance of greater than 75 TFlops from only CPU nodes in the past 5 years. IISc shall independently obtain inputs from the provided referees before arriving at a final decision.

11	Page 7 Section 6 Earnest Money Deposit (EMD)	EMD can be exempted for MSME bidders as per GFR. In this case, the MSME bidder should provide a valid MSME certificate from a competent authority in the technical bid.
12	Page 11 Section 11 Technical Bid – Terms and Conditions Item 11.i	i. Three references where they have carried out supply, installation and support of systems with peak performance of greater than 75 TFlops in the past 5 years. Modified to: i. Three references where they have carried out supply, installation and support of systems with peak performance of greater than 75 TFlops from only CPU nodes in the past 5 years.
13	Page 13 Section 14 Payment Terms Item number 5	Modified to: The total solution as per the agreed bill of materials must be supplied within 8 – 10 weeks after receiving a firm PO from IISc. The installation and acceptance must be completed within 2-3 weeks after supply of the equipment.
14	Page 13 Section 14 Payment terms	Addition of a new item, item 6: 6. Liquidated Damage: As time is the essence for this procurement, hence the ordered materials are required to be delivered and installed in all respects within the stipulated period in the purchase order failing which penalty for late delivery and installation will be imposed at the rate 1% of the total order value per week or part thereof for the delayed period subject to maximum of 10% of the total order value and this liquidated damage will be deducted during the payment of the invoice / bill of the supplier. Earliest / expected delivery period should be clearly indicated in the technical bid.
15	Page 14 Item 15 Important Dates	Modified to: 4. Release of corrigendum to the tender based on the queries, if necessary: February 28, 2020. 5. Start date for submission of the bid: March 2, 2020, 9 AM IST through online mode of CPPP. 6. Last date for submission of the bid: March 17, 2020, 6 PM IST through online mode of CPPP. 7. Opening of the technical bids: March 20, 2020, 10.30 AM IST through online mode of CPPP. 8. Technical presentation by the bidders: March 23, 2020. 9. Opening of the commercial bids: Will be intimated via CPPP.

2. Response to Queries

SNo	Reference	Query	Response
1	Page 5 Master node	We request the CPU on the master node alone be of a lower clock speed, which is, minimum 48 cores and minimum 2.4GHz. This is because of the TDP restriction on the system on account of the 2.6GHz or above CPU	Refer to SNo 3 in the corrigendum table above in this document.
2	Page 5 Master node	Another option on the master node could be to keep the specification same, but allow us to use a JBOD for the 500TB storage portion	The storage should be part of the master node. No external storage allowed.
3	Page 5 Master node Form Factor	The form factor asked is 1U/2U rack mountable chassis. We would like to understand if this will allow vendors to quote the dense form factor solution like the 4-nodes in a 2U chassis as well	Yes. Vendors are welcome to quote for the dense form factor solution.
4	General specs	In reference to the current RFP the overall compute only node translates to 36Nodes of AMD with Dual **7542** VS 30 Nodes of **6248R or 8268.** Intel is getting an undue advantage due to the way the RFP is written. ……. Query continued in the annexure after this table. **Recommendation to have Level playing field for both AMD and Intel:** 1. AMD Solution provides 2300+ Cores as	The technical specifications in the tender were arrived at and based on a comprehensive survey of the needs of the IISc HPC user community. The technical specifications in the tender wrt this query will stay unmodified.

		compared to 1440 Cores of Intel. … Points 2-3: Query continued in the annexure. **Requesting you to kindly reconsider our recommendation which will allow us to be more competitive in the entire solution.**	
5	Page 5 Master node Cross-compilation capability	Please note, there are very limited vendors who can support Nvidia Nvlink GPU on AMD based systems and the above clause restricts vendors to bid for an alternative system on GPU only nodes. Request you to please relax the above clause and allow bidders to quote for a heterogeneous solutions.	Refer to SNo 3 in the corrigendum table above in this document.
6	Page 7 Section 6 Earnest Money Deposit (EMD)	Request IISC to extend the NSIC exemption to the bidder - OEM / SI	Refer to SNo 11 in the corrigendum table above in this document.
7	Page 13 Section 14 Payment Terms Item 3	90% payment through LC and 10% payment post submission of bank gaurantee and installation. Request IISC to consider the same.	The payment terms specified in the tender wrt this query will stay unmodified.
8	Page 13 Section 14 Payment Terms Item number 5	Min. 8-10 weeks for Delivery is required & 4-6 weeks for installation from the date of delivery and site readiness. Request IISc to reconsider the same,	Refer to SNo 13 in the corrigendum table above in this document.
9	Page 5 Master node Form factor	Our query is for Achieving 500TB Usable after the RAID configuration we need to use multiple drives with form factor 1U & 2U it would be a difficult as we need to us multiple	Refer to SNo 3 in the corrigendum table above in this document.

		JBOD's , Hence request you to kindly allow only master node in 3U/4U Rackmountable Chassis.	
10	Page 4 1.a. CPU-only-Regular-Nodes SNo 5 NIC for MPI Communications and Page 6 2. Peripherals SNo 1 Interconnect switches and cables	In peripherals primary interconnect for MPI communication : its mentioned only infiniband switch with appropriate number of ports and cables to connect all the nodes mentioned above , Kindly confirm the interconnect speed of the switch.	Refer to SNo 1 and SNo 9 in the corrigendum table above in this document.
11	Page 5 1.b. Master Node Item 14 Storage	These features are specific to a particular vendor. We request that Vendors be allowed to quote external storage also we recommend to increase the cache from 2GB to 8GB.	No external storage will be allowed. The technical specifications in the tender wrt this query will stay unmodified.
12	Page 5 1.b. Master Node Item 14 Storage	We understand that the 500TB usable capacity should be in RAID6. Kindly confirm.	Yes.
13	Page 5 1.b. Master Node Item 14 Storage	Also kindly clarify the capacity and RAID level of the OS drives.	Capacity: Refer to SNo 4 in the corrigendum table above in this document. RAID level: Provide as appropriate.
14	Page 5 Master node Cross-compilation capability	We request to elaborate more on cross-compilation capabilities for creating executables for the CPU/GPU nodes.	The master node should have the capability for compilation to the corresponding binaries. Also, refer to SNo 5 in the corrigendum table above in this document.

15	Page 6 1.d. V100 Nodes	We request that the CPU architecture for GPU nodes be generalized as x86, as Nvidia with SXM2 NVLink is not available with AMD	The technical specifications in the tender wrt this query will stay unmodified.
16	Page 6 1.e. SSD Nodes	We request that this clause be relaxed to 1.6TB NVMe SSDs or more.	The technical specifications in the tender wrt this query will stay unmodified.
17	Page 6 2. Peripherals Item 1 Interconnect switches and cables	Kindly clarify if the primary interconnect needs to be a managed or externally managed switch	The bidder has the flexibility to provide the appropriate.
18	Page 8 Section 8 Acceptance Criteria Item 1	Kindly clarify whether the Tflops of Head node will be included in calculation of sustained 85TFlops or it has to come exclusively from the CPU-Only Regular Compute Nodes.	The performance of all CPU-Only Regular nodes including the compute and master nodes will be included in the calculation of sustained 85 TFlops acceptance criteria.
19	Page 7 Section 5 Bidder's Eligibility Criteria and Page 11 Section 11 Technical Bid – Terms and Conditions Item 11.i	We request that this clause be relaxed to "Three references where OEM/Bidder have carried out supply, installation and support of systems with peak performance of greater than 75 TFlops in the past 5 years.	The terms specified in the tender wrt this query will stay unmodified. Also, refer to SNo 10 and SNo 12 in the corrigendum table above in this document.
20	Page 4 and Page 5 CPU and GPU nodes	Can the CPU and GPU nodes have different processor OEMs?	Yes.
21	Memory – All nodes	What is the expectation in terms of balanced ? Do we need to	Yes.

		populate all the memory channels in the Server ?	
22	Storage - CPU, GPU, High Memory and Nodes with SSDs Diskless boot	Kindly make it optional and allow vendors to opt for a disk if needed.	The terms specified in the tender wrt this query will stay unmodified.
23	Page 4 - Form Factor	Is the requirement for a Chassis with dense nodes OR a Server with 1U or 2U Form factor. Kindly leave the choice on the vendor.	Refer to SNo 2 in the corrigendum table above in this document.

ANNEXURE

Query related to AMD vs Intel Comparison

In reference to the current RFP the overall compute only node translates to 36Nodes of AMD with Dual **7542** VS 30 Nodes of **6248R or 8268.**
Intel is getting an undue advantage due to the way the RFP is written.

Here are few points below on why AMD should be your choice of processors for HPC:

1. Memory bandwidth: AMD offers you the industry best memory bandwidth by supporting 8 memory channels and 3200Mhz of memory frequency. This translates to a theoretical memory bandwidth of 409GB/s as against Intel's 281GB/s.
2. More system I/O: AMD supports you with 128 PCIe Gen4 lanes on Rome platform against 48 PCIe Lanes of Intel.
3. Industry best spec numbers on spec.org
4. First 7nm process technology which ensures your TDPs are in control.
5. We do support processors up till 64cores per processor.

		Intel® Xeon® Gold 8268 Processor	**AMD EPYC 7542**
1	# of CPU Cores	24 Cores	32 Cores
2	# of Threads	40 Threads	64 Threads
3	Base Frequency	2.90Ghz	2.90Ghz
2	Size of L1/L2 Cache/L3 Cache	35.75MB(Smart Cache)	128MB of L3 Cache
3	FLOPS Theoretical	4.4544	2.9696
4	HPL score & other BM scores like Rpeak, Rmax & SPEC	HPL Efficiency: 65%	HPL Efficiency: 87%

		SpecInt base: 316 Specfp base: 267	SpecInt base: 394 Specfp base: 345
6	Max Memory capacity per proc	768GB using DDR4	4TB using LRDIMMs
7	Memory Bandwidth per processor	140.7GB/s	204.8GB/s
8	AVX Vector width	AVX512(512Bit). But at least 95% of the applications do not utilize the AVX512 and using it may impact the base frequency. Attaching the AVX frequencies in Intel.	AVX2(256Bit)

Recommendation to have Level playing field for both AMD and Intel:

1. AMD Solution provides 2300+ Cores as compared to 1440 Cores of Intel. Request you to please incorporate the minimum requirement to 2300Cores which will bring both the solutions on par.
2. AMD Solution has to be provided with 18TB of memory as compared to 11TB of memory on Intel. Request you to please make 18TB of memory as base line for all solutions.
3. AMD solution delivers a system level memory bandwidth of 14.4TB/s as compared to 8.4TB.s. Memory bandwidth is the key for any HPC application performance and request you to please incorporate the requirement of 14TB/s system level memory bandwidth in the RFP.

Requesting you to kindly reconsider our recommendation which will allow us to be more competitive in the entire solution.