

	<p style="text-align: center;">भारतीय प्रौद्योगिकी संस्थान मद्रास चेन्नै 600 036  <b>INDIAN INSTITUTE OF TECHNOLOGY MADRAS Chennai 600 036</b>  भंडार एवं क्रय अनुभाग  <b>STORES &amp; PURCHASE SECTION</b>  दूरभाष: (044) 2257 8285 / 8286 / 8287 / 8288 फ़ैक्स: (044) 2257 8292 / 2257 8082  Telephone : (044) 2257 8287, 8285, 8288  email ID-<a href="mailto:adstores@iitm.ac.in">adstores@iitm.ac.in</a>  GSTIN : 33AAAAI3615G1Z6</p>	
---	--	---

**G. Chitrapavai**  
**Deputy Registrar (Stores & Purchase)**

**Dated : 05.01.19**

### **CORRIGENDUM**

**Tender No. IITM/SPS/CC/HPCE Cluster/008/2018-19**

**Tender ID: 2018\_IITM\_19977\_1**

**“Supply and Installation of HPCE Cluster“**

In the tender document,

- a) In Page No. 12 (Annexure A) and in Page No. 20 (Annexure B) – Point No. 1 to 3 may be read as**

- |   |
|---|
| <ol style="list-style-type: none"> <li>1. <b>OEM</b> must have at least 5 different entries in latest Top 500 supercomputers listed at top500.org (necessary proof of document to be enclosed with tender).</li> <li>2. <b>OEM</b> must have presence for at least 10 years in this business. (Proof of selling high performance computing facilities for past 10 years should be attached).</li> <li>3. <b>OEM</b> must have installed <b>at least 1 X 100 TF system (CPU only)</b> in India with at least 100 TB of parallel file system. Order copy and work completion certificate from the client should be attached.</li> </ol> |
|---|

#### **Instead of**

- |   |
|---|
| <ol style="list-style-type: none"> <li>1. Vendor must have presence in Top 500 supercomputers listed at top500.org. (Necessary proof of document to be enclosed with tender).</li> <li>2. Vendors must have presence for at least 10 years in this business. (Proof of selling high performance computing facilities for past 10 years should be attached).</li> <li>3. Vendor must have installed at least 5 X 200 TF system (CPU only) in India with at least 200 TB of parallel file system. Order copy and work completion certificate from the client should be attached.</li> </ol> |
|---|

- b) In Page No.13, the below point may be added as**

- |   |
|---|
| <ol style="list-style-type: none"> <li>i) The interconnect &amp; network should provide switches, cables and ports to accept 125% of the nodes on day 1.</li> </ol> |
|---|

c) In Page Nos. 14, 21, 23-24, 25-26 & 27-28 – Point Nos. 1, 2, & 4 may be read as

S.No.	Item	Technical Specification
1	CPU-only compute nodes	<ul style="list-style-type: none"> <li>- 500/600/700/800 TFlops of compute power in double precision.</li> <li>- clock at least 32 instructions per core per cycle.</li> <li>- at least 16 cores per socket with 2.4GHz or more.</li> <li>- At least 192 GB DDR4 RAM for main memory in a balanced configuration per node with 2666 MHz.</li> <li>- At least two processors per node.</li> <li>- <b>At least 2 TB 7.2K RPM SATA Hard disk per node.</b></li> <li>- Rack mountable with suitable mounting kit.</li> <li>- Redundant power supplies for all nodes.</li> <li>- Service level: Next business day.</li> </ul>
2	CPU-GPU compute nodes	<ul style="list-style-type: none"> <li>- 200 TFlops of compute power in double precision; Two GPUs per node</li> <li>- <b>NVIDIA Tesla V100 32 GB cards.</b></li> <li>- CPU: x86-64 architecture, at least 2.4 GHz clock.</li> <li>- at least 32 instructions per core per cycle, at least 16 cores per socket.</li> <li>- At least 192 GB main memory in a balanced configuration per node with 2666 MHz or more.</li> <li>- <b>At least 2 TB 7.2K RPM SATA Hard disk per node.</b></li> <li>- Redundant power supplies.</li> <li>- Rack mountable with suitable mounting kit.</li> <li>- Service level: Next business day.</li> <li>- Vendor should quote for PCI-e and NVLink separately, if applicable.</li> </ul>
4	<b>Service nodes (Two Master, One login, I/O as required by vendor design)</b>	<ul style="list-style-type: none"> <li>- x86-64 architecture, at least 2.4 GHz At least 192 GB main memory in a balanced configuration.</li> <li>- <b>At least 2 TB SATA 7.2K RPM Hard disk per node.</b></li> <li>- Number of nodes required for the optimal parallel file system with redundancy.</li> <li>- <b>Master nodes should be in High Availability mode. Job scheduler, cluster management and monitoring tools should be available in failover capability.</b></li> <li>- <b>Master Nodes and Login Node to be provided with two or more network port having 10Gbps, SFP+ compatible to single mode fiber to connect with the campus LAN.</b></li> <li>- <b>Master Nodes and Login Node should have dedicated management port and two or more 1 Gbps Ethernet ports and infiniband ports as required for the optimized solution.</b></li> <li>- Rack mountable with suitable mounting kit.</li> <li>- Service level: 24x7.</li> </ul>

Instead of

S.No.	Item	Technical Specification
	CPU-only compute nodes	<ul style="list-style-type: none"> <li>- 500/600/700/800 TFlops of compute power in double precision</li> <li>- clock at least 32 instructions per core per cycle.</li> <li>- at least 16 cores per socket with 2.4GHz or more.</li> <li>- At least 192 GB DDR4 RAM for main memory in a balanced configuration per node with 2666 MHz.</li> <li>- At least two processors per node.</li> <li>- At least 2 TB 15K RPM SAS Hard disk per node.</li> <li>- One free PCI-e slot for future expansion.</li> <li>- Rack mountable with suitable mounting kit.</li> <li>- Redundant power supplies for all nodes.</li> <li>- Service level: Next business day.</li> </ul>

2	CPU-GPU compute nodes	<ul style="list-style-type: none"> <li>- 200 TFlops of compute power in double precision; Two GPUs per node.</li> <li>- NVIDIA Tesla V100 32 GB cards with at least 16 GB main memory.</li> <li>- CPU: x86-64 architecture, at least 2.4 GHz clock.</li> <li>- at least 32 instructions per core per cycle, at least 16 cores per socket.</li> <li>- At least 192 GB main memory in a balanced configuration per node with 2666 MHz or more.</li> <li>- At least 2 TB SAS 15K RPM Hard disk per node.</li> <li>- Redundant power supplies.</li> <li>- Rack mountable with suitable mounting kit.</li> <li>- Service level: Next business day.</li> <li>- Vendor should quote for PCI-e and NVLink separately, if applicable.</li> </ul>
4	Service nodes (Master, login, I/O, etc.)	<ul style="list-style-type: none"> <li>- x86-64 architecture, at least 2.4 GHz At least 192 GB main memory in a balanced configuration.</li> <li>- At least 2 TB SAS 15K RPM Hard disk per node.</li> <li>- Number of nodes required for the optimal parallel file system with redundancy.</li> <li>- Rack mountable with suitable mounting kit.</li> <li>- Service level: 24x7.</li> </ul>

**d) In Page Nos .15, 22, 24, 26, & 28 – Point Nos. 6, 8, 9 &10 may be added as**

6	Network/ Interconnect	<ul style="list-style-type: none"> <li>- At least 100 Gbps per port bandwidth.</li> <li>- <b>Aggregate link speed of 100Gbps.</b></li> <li>- <b>MPI latency of &lt;= 1.5 micro seconds for 8 byte messages with error correction and detection always enabled.</b></li> <li>-<b>All switches should have redundant power supplies.</b></li> <li>- <b>Managed Gigabit Ethernet switches to connect all HPC nodes for cluster management.</b></li> </ul>
8	Cooling	<ul style="list-style-type: none"> <li>-As required for the solution above using water or air.</li> <li>-The cooling should be for 125% of the calculated load.</li> <li>-The vendor is expected to visit the site to understand the existing infrastructure.</li> <li>-<b>Redundancy of cooling system should be N+1 at all points.</b></li> </ul>
9	Operating system/ Cluster Management	Licensed OS(64-bit Linux or Unix variant) and licensed Cluster management / monitoring software should be fully supported (by the OEM), with updates during warranty period <b>with support for the full 125% of nodes.</b>
10	Compilers, libraries and tools	Vendor must supply <b>10 Concurrent</b> Licenses of Intel Parallel Cluster Studio which should be valid for the entire warranty period (5 years).

**Instead of**

6	Interconnect	- At least 100 Gbps per port bandwidth. Adequate redundancy to avoid single-point failure in the interconnect should be provided. All switches should have redundant power supplies.
8	Cooling	As required for the solution above using chilled water or air;The cooling should be for 125% of the calculated load. The vendor is expected to visit the site to understand the existing infrastructure.

9	Operating system/ Cluster Management	Licensed OS(64-bit Linux or Unix variant) and licensed cluster management/monitoring software should be fully supported (by the OEM), with updates during warranty period. Cluster management software should be licensed without the limits on the number of nodes.
10	Compilers , libraries and tools	Support for the complete software suite including all the software currently supported on the Virgo cluster at IITM (Abaqus, Comsol, Mathematica, Gaussian, Nastran, Dytran, Marc, Ansys/Fluent, NAMD, LAMMPS, GROMACS, Amber, Accelrys, Matlab.) along with various compilers (such as GNU GCC collection, Intel, Java compiler.). Vendor must supply 10 Licenses of Intel Parallel Cluster Studio which should be valid for the entire warranty period (5 years).

e) In Page Nos.16, 22, 25, 27 & 29 - Point No. 13, may be read as

13	Hardware reliability	- Redundant power supplies and voltage regulator modules.
Instead of		
13	Hardware reliability	- Redundant paths for all system RAID. - Redundant power supplies and voltage regulator modules.

f) In Page No.17 - Point No. 5 (iii & iv), may be read as

iii) The Servers should be factory tested and burn-in reports provided.		
iv) Installation of the HPC Cluster hardware, OS, software components must be done by trained engineers from the OEM. The solution should be validated and certified by the OEM.		
Instead of		
iii) The cluster must be factory integrated and brought to IIT Madras.		
iv) Installation of the factory-integrated hardware, OS, software components must be done by trained engineers from the solution provider. The solution should be validated and certified by the provider.		

g) In Page No.19 - Point Nos. 8 & 11 may be read as

8) <b>Support:</b> Submit the specifics of your 24x7 support in terms of response and resolution time for various types of problems in the HPC cluster. <b>OEM should give a SPOC(Single Point Of Contact) to support the overall solution provided.</b>		
11) The space currently available should be utilized for the erection of the HPCE and we will not provide additional space for the installation of this new equipment which is proposed to be purchased. <b>Dimensions of Data Centre to host the solution provided by the vendor, height &amp; width of the doors and load bearing capacity is given in Annexure-E.</b>		
Instead of		
8) <b>Support:</b> Submit the specifics of your 24x7 support in terms of response and resolution time for various types of problems in the HPC cluster.		
11) The space currently available should be utilized for the erection of the HPCE and we will not provide additional space for the installation of this new equipment which is proposed to be purchased.		

h) In Page No.30 (BOQ Price Bid format, Annexure - C ) – S. Nos. 3 & 11 may be read as

<b>3) NVIDIA Tesla V100 32 GB cards</b> <b>11) Cooling System (Inclusive of Tower if required).</b>
<b>Instead of</b>
3) NVIDIA Tesla V100 32 GB cards with at least 16 GB main memory 11) Cooling – Tower Cooling System.

i) The Annexure - E added newly.

j) The last date, due date for receipt of tender mentioned in the tender document may be read as

**25.01.2019 before 02.00 p.m.**

k) The date & time of opening of tender mentioned in the tender document may be read as


**28.01.2019 before 03.00 p.m.**

All other conditions remains the same.

**Sd/-**  
**Deputy Registrar**  
**Stores & Purchase**

# ANNEXURE-E

## LAYOUT



- SERVER AREA dimension L X B X H – 11.9 m X 11.6 m X 2.75 m
- Available Area for New HPC Cluster – 11.9 m X 5.8 m X 2.75 m
- False floor depth – 0.6 m
- Load bearing capacity of floor tiles – 2200 kg.