



**INDIAN INSTITUTE OF TECHNOLOGY BOMBAY**

**MATERIALS MANAGEMENT DIVISION**

**Powai, Mumbai 400076.**

**Item Description : GPU Server with 2x EPYC 9754 and 4x H200**

**PR No.1000050021**

**(Rfx No.6100002388)**

<b>Sr. No</b>	<b>Detailed Technical Specification</b>	<b>Technical Compliance (Yes / No)</b>	<b>Additional Information ( if any )</b>
1.	Processor : AMD EPYC 9754, 2.25Ghz, 256MB L3 Cache, 128C/256T, 360W or Equivalent (QTY : 02 Nos.)		
2.	Barebone Server : 4U Rackmount with 3+1, 80-PLUS Platinum, 2000W CRPS (QTY : 01 No.)  Dual Socket SP5 (LGA 6096), supports AMD EPYC™ 9005*/9004 (with AMD 3D V-Cache™ Technology) and 97x4 series processors  12+12 DIMM slots (1DPC), supports DDR5 RDIMM, RDIMM-3DS  16 hot-swap 2.5" NVMe (PCIe5.0 x4) drive bays or 24 Hot-swap 2.5" SATA/SAS* drive bays *Additional RAID/HBA card required  10 FHFL dual-slot PCIe5.0 x16 + 2 FHFL single-slot PCIe5.0 x8 shared space with adjacent dual-slot (with PEX89104 PCIe Switch)  1 FHFL single-slot PCIe5.0 x16 (CPU direct link), 2 FHHL PCIe5.0 x16  Supports 2 M.2 (PCIe3.0 x4 or SATA 6Gb/s)  1 OCP NIC 3.0 (PCIe5.0 x16), 2 RJ45 (1GbE) by Intel® i350 Remote Management (IPMI)		

i.	Memory : 128GB DDR5 ECC RDIMM	8	
ii.	Storage (OS) : 960GB U.2 NVME Drive (Raid 1)	2	
iii.	Storage (Data) : 7.68TB Enterprise SATA SSD (Raid 5)	6	
iv.	Raid Controller : 16 Port Internal Raid Card 12G Trimode 8 GB Dual Slim Line SAS X8 - 9560-16i	1	
v.	Cables : SAS Cables	2	
vi.	LAN : ConnectX@-6 Dx EN Network Interface Card, 100GbE Dual-Port QSFP56, PCIe4.0 x 16 (SFP Modules are not included)	1	
vii.	GPU : Nvidia H200(141GB) NVL GPU Card or equivalent	4	
viii.	Processor :Two processors per system with each socket having the following:		
ix.	Processor has to be at least PCIe gen 5.0 compliant.		
x.	Processor Physical Cores should be minimum 128 with 256 Threads with minimum 250MB Cache		
xi.	At least 2.25 GHz base frequency per core		
xii.	TDP per socket must be within 400 W		
xiii.	Should support AVX256 and AVX2 instruction sets		
xiv.	Must support virtualization / hypervisor		
xv.	Memory : 24 Memory DIMM Slots Supporting DDR5 Server Memory		
xvi.	Total minimum 1024GB ECC Registered DDR5 memory with system running frequency of 4800 MHz		
xvii.	Every memory channel from the processor(s) should be populated in a balanced configuration.		
xviii.	Storage : Each server should have 16 hot-swap 2.5" NVMe (PCIe5.0 x4) drive bays or 24 Hot-swap 2.5" SATA/SAS* drive bays		
xix.	At least 2 x 960GB Enterprise 2.5" U.2/U.3 NVME to be populated in Raid 1 configuration from day one		
xx.	Hotswap bays to be populated with 6 Nos. 7.68 TB SATA SSD(Generation-4) capacity in Raid 5 via 16Port		

	Hardware Raid Controller, supporting Trimode and minimum cache of 8GB		
xxi.	Network : On-board 2 Gigabit RJ45 Ports or higher + 1 No. of Dual Port 100GbE QSFP56 or Higher		
xxii.	PCI expansion Free : 5 Nos. PCIe Gen5 X16 Slots and 2 Nos. PCIe Gen5 X8 Slots should be free for future expansion		
xxiii.	External Ports : Following minimum number of external ports should be available in the system		
xxiv.	1 x VGA ports		
xxv.	2 x USB 3.2 ports		
xxvi.	1 x dedicated 1GbE RJ45 port for IPMI		
xxvii.	Co- Processor/ Accelerator : Each server has to be populated with 8 numbers of accelerators. Each accelerator should have:		
xxviii.	1. 141GB or more GDDR6 with ECC		
xxix.	2. 4.8TB/second or higher memory bandwidth		
xxx.	3. 16xPCIe Gen5 interface interconnect each capable of 128GB/s or higher bidirectional throughput		
xxxi.	4. 16896 accelerator cores or more		
xxxii.	5. 60 TFLOPS FP32 (peak) performance		
xxxiii.	7. Max power consumption of 700 W		
xxxiv.	8. Passive cooling support		
xxxv.	9. Virtual accelerator support		
xxxvi.	10. Should have 900GB/s additional Communication Link other than PCIe Gen5 : 128GB/s between all accelerators.		
xxvii.	Chassis : Appropriate Chassis (Maximum 4U) with rail-kit. The network, display and USB ports have to be on the front and rear side. Supplied chassis has to be approved by the system board OEM. Chassis should be populated with a maximum number of redundant cooling fans that can be accommodated in the chassis in the cumulative air flow should be sufficient to take away the heat generated at peak load with fully populated hardware.		

xxviii.	Server Management : Dedicated IPMI 2.0 compliant management 1 Gbps ethernet port having support for		
xxxix.	1. System health monitoring		
xl.	2. PMBus and SMBus monitoring for redundant power supply and storage backplane		
xli.	3. Event log accounting and monitoring changes in the server hardware and system configuration		
xlvi.	4. Virtual media over network and Virtual KVM (KVM over IP).		
xliii.	5. Agentless management using the out - of - band remote management port		
xliv.	6. At least 128 bit SSL encryption and secure shell Version 2 support.		
xliv.	7. Should provide support for AES and 3DES on browsers.		
xlvi.	8. Should provide remote firmware update functionality.		
xlvii.	9. Should provide support for Java free graphical remote console.		
xlviii.	10. IPMI must have graphical interfaces accessible to qualified users only and through browsers and give access to all of the above.		
xlix.	11. All required licenses (if any) to use IPMI features		
i.	Dedicated IPMI 2.0 compliant management 1 Gbps ethernet port having support for		
ii.	80+ Platinum Certified Redundant Power Supply.		
iii.	Power supplies should be only of the make/model approved for the motherboard. Power supply has to be configured in N+1 redundancy.		
liii.	The system should sustain in the occasion of failure of one power supply. Power Management Bus (PMBus) support should be provided and status information must be available via IPMI.		
liv.	Certification : The server and all its components should be verified and recommended by the system board manufacturer by means of compatibility list. The specifications of critical components such as		

	CPU, accelerator, memory and disks should be verifiable from the websites of respective manufacturers.		
iv.	Power Cord and other cables : Power cable should have IEC 320 (male-female) connectors matching the peak-load rating of the power supplies. All the network cables required are to be supplied in appropriate numbers and of appropriate specification.		
lvi.	OS Support : The system should support the following minimum operating systems without any deviations:		
lvii.	Windows Server 2019		
lviii.	Windows Server 2022		
lix.	Red Hat Enterprise Linux server 8.6, 8.7, 8.8, 9.0, 9.1 & 9.2 x64		
lx.	SUSE Linux Enterprise server 15 SP4 & 15 SP5		
lxi.	Ubuntu 20.04.5 LTS x64, 20.04.6 LTS x64, 22.04 LTS x64, 22.04.1 LTS x64, 22.04.2 LTS x64 & 22.04.3 LTS x64		
lxii.	VMware ESXi 7.0 Update 3i, ESXi 8.0, ESXi 8.0 Update 1, Citrix Hypervisor 8.2 LTSR CU1		
lxiii.	Additional Requirements : The main-board of the system should be from the same brand of the supplied entire system. The main-board PCB should be " permanently printed" with the system brand markings but not through pasted stickers or removable markings. Evidence towards the same needs to be submitted along with the bid.		
lxiv.	Warranty : 3 years comprehensive on-site advance replacement warranty. The warranty should be trackable online on the OEM website through system serial number providing detailed system hardware configuration report.		
lxv.	Environment compliance : RoHS compliance		
lxvi.	Data-Sheet : It is mandatory to submit the technical data-sheet of the offered system along with the bid. The datasheet should match the technical compliance documents of bid, failing which the bid shall be disqualified without any query.		

lxvii.	MII (Make in India) : It is mandatory to submit declaration of breakup of MII(Make in India) content with proof.		
lxviii.	MAC Address : The MAC address OUI (Organizationally Unique Identifier) of the servers, and any other hardware required to make the system complete must be registered in the name of quoted OEM of the product.  Self-declaration by OEM on their letter head signed by the authorized signatory. The same will be checked on: <a href="https://regauth.standards.ieee.org/standards-raweb/pub/view.html#registries">https://regauth.standards.ieee.org/standards-raweb/pub/view.html#registries</a>		
lxix.	System Utility with AI Frameworks Pre-Loaded : System should be preloaded with Precompiled frameworks (CPU & GPU optimized MxNet, CuDDN, Caffe and Pytorch) to be supplied with the system, license must be in the name of customer organization.		
lxx.	Job Scheduler : Datasheet of the Job Scheduler is compulsory to submit along with the bid.		
lxxi.	MAF & BOM : The Vendor must provide Manufacturer Authorization Form (MAF) & technical Bill Of Material (BOM) with the quote.		
lxxii.	Submit Test Reports for Validation of Offered or Technical Evaluation of Item.		
lxxiii.	The Bidder will be responsible of hardware Installation within 2 weeks of supply.		
lxxiv.	Bidders with past experience of dealing with IITs, NITs or Research Institutions in India and having their service centre in Mumbai are preferable		
lxxv.	Bidders should submit list of engineers who will be available for Instant Support . Knowledge of Linux is essential and please highlight your team accordingly.		
lxxvi.	Bidder should submit ISO certifications along with other OEM Certifications.		