



INDIAN INSTITUTE OF SCIENCE EDUCATION AND RESEARCH PUNE

PREBID CLARIFICATION ON GeM Bid NUMBER - GEM/2022/B/2331134

Item Description- Procurement of High Performance Computing Cluster.

Refer GeM Bid published on Government e Marketplace (GeM) and Institute website www.iiserpune.ac.in Dt. 11/07/2022. The press tender notice was issued in Indian Express Mumbai, Pune & Delhi Editions.

Pre-Bid meeting was held on 20/07/2022 at 3.00 PM and minutes of meeting is as under:

At the outset, the Chairman welcomed all the Members and the representative of the Prospective Bidders and briefed in general the scope of the tender and thereafter requested Assistant Registrar (S&P) to brief the bidders on the salient features of the tender.

The representatives present were satisfied with the replies given and it was informed that the corrections / additions / clarifications given, as discussed during the Pre-Bid Conference would be hosted on the website of IISER Pune and all the Prospective Bidders are required to take cognizance of the proceedings of the Pre-Bid Conference before submitting their bids as stipulated in the Bidding Documents.

**The last date for submission of tender due date is extended up to 22 Aug 2022
- Time - 3.00 pm**

The other terms & conditions of the notice issued on our IISER website www.iiserpune.ac.in will remain unchanged. No more correspondence in this regard will be entertained.

The meeting ended with vote of thanks to the Chair

20/07/2022

Sd-x-x-x
Assistant Registrar (S&P)

TECHNICAL AND COMMERCIAL QUERIES AND CLARIFICATION

PRE-BID CONFERENCE FOR PROCUREMENT OF HIGH PERFORMANCE COMPUTING CLUSTER

S. No	Technical specification clause No.	Query/Clarification Sought	Clarification / Amendment
1.	H. 1	Request to relax number of CPU cores to 6 from minimum required 8.	Accepted. The specification is now amended to: "One Processor with minimum <u>6</u> Cores and minimum base clock frequency of 2.5 GHz <u>for the latest generation of CPU from the OEM</u> ".
2	H. 6	Request to amend the requirement of chassis size from Mid-Tower to Tower	Accepted.
3.	H. 7	Request to drop VGA port as a mandatory requirement	Accepted. The requirement of <u>VGA</u> port is dropped from the minimum specification requirement.
4.	E. 3	Clarification requested: whether a hardware RAID controller is required to be provisioned for data protection?	Yes the Hardware raid controller is required. The specification is now amended to be read as below: "5 x 12TB NLSAS 7.2K RPM 3.5" HDD 2 x 960GB SATA SSD <u>With hardware raid controller minimum supporting RAID 0,1,10 & 5 disk configurations.</u> "
5.	E. 5	Request for removing GPU capability feature in terms of the expansion slots as required with the server.	Accepted. The specification is now amended to be read as below: " Min. 4 x PCIe Low-Profile Slots 4.0 x8 atleast"
6.	E. 7	Request to lower down 'Network interface' specification in terms of requirement of '4 x Dual port 25G SFP+' ports.	Accepted. The specification is now amended to be read as below: "1x 1 GbE Gigabit Ethernet LAN port, 2x 10G/1G Base-T Ethernet LAN ports 1 x 100G HDR HBA, <u>Dual port 25G SFP+</u> "
7.	I. 1	Request for relaxing stackability requirement for the switch	Accepted. The specification is now amended to be read as below: " 48 x 10/100/1000BASE-T ports."
8.	G. Title	Request to change the title of the equipment	Accepted. The equipment title / name is changed to ' Infini-band

		' Infini-band Switch HDR 100' to 'Infini-band Switch HDR 200'	Switch HDR 200 for Primary Network'
9.	A. 9	Request to relax the mandatory requirement of at least 2 USB 3.0 to be present on the rear side of the server	Accepted. The specification is now amended to be read as below: ' 2x USB 3.0 or more, 1x VGA'
10	C. 1	Clarification request: whether single port consumer grade NVMe SSDs or dual port data center grade NVMe SSDs are required?	Clarification: The bidders are requested to quote for Data Center Grade or Enterprise class NVMe SSDs in their solution. The solution thus offered should meet all performance metrics as detailed in the Technical specifications.
11	C. 1	Clarification request: whether bidders can quote SAS attached SSDs to meet Required capacity and performance?	No change in the technical specifications. Bidders are requested to quote for solution with NVMe SSDs (Data Center Grade or Enterprise class) only.
12	C. 2 (Bullet Serial No. 05)	Request for amendment to bind the specifications to support latest versions of GPU Direct Storage (GDS)	Accepted. The specification is now amended to be read as below: 'The solution must be POSIX compliant and support NFS, SMB , S3 protocols and latest versions of GPU Direct storage(GDS)'
13	C. 4	Request to confirm / include the capacity for metadata in terms of number of files to be accommodated with the technical specifications	Accepted. The solution should be capable of holding metadata with a capacity to accommodate 1 Billion files.
14	C. 5 (Bullet Serial No. 02)	Confirm requested: Whether single CPU based controllers/nodes are acceptable?	It is clarified that this aspect is left to the liberty of the solution design architect.
15	C. 5 (Bullet Serial No. 03)	Clarification request: Confirmation upon RAID configuration / Layout.	It is clarified that the bidders are required to quote for the solution with RAID 6 (8+2) disk configuration.
16	C. 5 (Bullet Serial No. 07)	Clarification request: Specification for the maximum permissible disk rebuilding time	It is clarified that the maximum permissible rebuilding time for disk(s) allowed is 12 Hours upon replacement.
17	C. 6 (Bullet Serial No. 01)	Request for amendment: In context with IOR or FIO benchmarks/Industry standard benchmarks; the specification is requested to be amended to '1MB or in the multiples thereof' instead of 1 MB as required.	Accepted. The specification is now amended to be read as below: 'Running on compute nodes with 1 MB transfer size or in the multiples thereof and file size double than total storage cache and I/O node memory. The total transfer size should be greater than the RAM available on the AI compute nodes.'
18	C. 6 (Bullet Serial Nos. 02 & 03)	Clarification request: Whether all or only certain number of the cores of clients nodes can be used to demonstrate the performance required.	It is clarified that this aspect is left to the liberty of the solution design architect as long as the I/O performance benchmark can be practically demonstrated with the entire set of solution.

19	C. 6 (Bullet Serial No. 04)	Request to amendment to include : 4K read performance of 3 million IOPS or higher with DIRECTIO	No change in the technical specifications.
20	D. 1	Request to confirm if the 7 GBps throughput has to be measured from Infiniband or 10/25/40 Gbps LAN network.	It is clarified that the performance should be measurable over LAN interfaces only
21	D. 1	Request to confirm / include the capacity for metadata in terms of number of files to be accommodated with the technical specifications.	Accepted. The solution should be capable of holding metadata with a capacity to accommodate 2 Billion files.
22	A. 3	Request to lower down the requirement pertaining to number of prepopulated GPU cards from existing 8 to 4	No change in the technical specifications.
23	A. 4	Request to lower down the requirement pertaining to number of prepopulated U.2 NVMe drives from existing 4 to 2	No change in the technical specifications.
24	A. 5	Request to lower down the requirement pertaining to number of Single Port NVIDIA ConnectX-6 200Gb/s InfiniBand Ports from existing 4 to 2	No change in the technical specifications.
25	A. 10, B. 14	Request to consider / include IEC 19 type Power cable compatibility to the power supplies	Accepted. Bidders to note that compatible power plugs sockets over the rack power strips are required to be factored in.
26	A. 12, B. 7 & E. 6	<p>Request to include the following in the specification:</p> <ul style="list-style-type: none"> IPMI 2.0 or equivalent Support with KVM and Media over LAN features. Must include any licenses, if required for using these features. It should be able to automate mgmt. tasks and automated firmware updates. <p>System Management Security Features:</p> <ul style="list-style-type: none"> System shall support immutable hardware silicon root of trust to authenticate firmware booting and updates and protect hardware from malicious code 	No change in the technical specifications.

		<ul style="list-style-type: none"> • System shall support digitally signed secure firmware updates • System shall support Automatic BIOS and Firmware Recovery from malicious changes • System should support instant secure erasure of drives. • System should support system Lockdown mode • System should support FIPS/Common criteria compliance • System should support UEFI secure boot with custom certificates 	
27	B. 5	Request to lower down the requirement pertaining to number of HDD bays from existing 12 to 8	Accepted. The specification is now amended to be read as below: HDD bays: ' <u>8</u> x 3.5" or 2.5" hot-swap drive bay or more.'
28	B. 10 & E. 8	Request to relax the mandatory requirement of at least 2 USB 3.0 to be present on the rear side of the server	Accepted. The specification is now amended to be read as below: '2x USB 3.0 or more, 1x VGA, 1x Management'
29	F. 1 (Bullet Serial No. 01)	Request to lower down the switch configuration / requirement	<p>The total quantity of 12 network switches as specified in the tender document are now grouped between two set of configuration / specifications</p> <p>The specification is amended to:</p> <p>Two switches (Qty-02) with At least 40 x 100/1000 BASE-T ports & 8 mGig ports up to 10G (Total 48 - RJ45 ports). 2x 1/10/25G fixed uplinks over SFP28, Minimum (non-stack) standalone switching capacity 340 Gbps, forwarding rate 250 Mpps, 1U Rack mountable.</p> <p>Ten switches (Qty-10) with At least 48 x 100/1000 BASE-T ports (Total 48 - RJ45 ports). 4 x 1/10 G uplinks over SFP/SFP+ transceivers, Minimum (non-stack) standalone switching capacity 150 Gbps, forwarding rate 120 Mpps, 1U Rack mountable.</p>
30	F. 1 (Bullet Serial No. 07)	Request for clarification: Exact number of optical fibre transceivers required per Network Switch.	It is clarified that total number of optical fibre transceivers required to be supplied (Not to be calculated on per network switch basis) is: 2 Pairs of each (1 pair of Single mode and 1 of Multimode transceivers) for 1G/10G SFP+ & 10G/25G SFP28 Cisco transceiver modules: Total 4 Pairs (2 S.M. + 2 M.M.) equated to 8 transceiver modules.

31	F. 3	Request for clarification: The exact number of UTP patch cords to be included with the supply.	It is clarified that the successful bidder has to additionally supply 800 Factory molded Cat 6 UTP Patch cords (2 Meter in length) over and above the quantity of those required for successful installation, testing and commissioning of the entire Setup.
32	A. 10	Request to relax the Power Supply ratings from 'Titanium' to 'Platinum' or better	Accepted. The specification is relaxed to accommodate offers with 'Platinum or better' as power supply ratings.
33	B.	Reduction in the quantity of required 'AI based Compute Nodes'	It had been decided by the competent authority to lower the quantity of the required 'AI based Compute Nodes' from existing <u>28</u> to <u>20</u> .

Note: The revised technical specifications and scope of work document had been compiled factoring-in the above discussed changes / amendments to the tender specifications (refer following Annexure - III). The bidders are advised to treat this document as final and are requested to comply accordingly.

Annexure-III

A) High End Accelerated Compute Node

Quantity		2 Nos.
Sr. No.	Specification	Description
1.	Processor	Dual x86_64 Processor CPU with minimum 64 cores/128 threads minimum base clock frequency of CPU 2.25 GHz, Support for Virtualization technology.
2.	Memory	Total 2TB Memory DDR 4, Min 3200 MHz or better with Advanced ECC with multi-bit error protection
3.	GPU	8x NVIDIA A100 80GB Tensor Core GPUs with Nvlink (GPU Memory 640GB total)
4.	Storage	OS: 1 x 1.92TB SATA SSD drives Internal Storage: 30TB (4 x 7.6 TB) U.2 NVMe drives
5.	Network interface	4 x Single Port NVIDIA ConnectX-6 200Gb/s InfiniBand. At least 2x10 G RJ45 Ethernet port network adapter.
6.	Operating System	Support latest version of 64-bit RHEL/CentOS/Ubuntu/Open source Linux – latest stable releases
7.	Form Factor	6U or less, Rack mountable, Rail Kits to be included - compatible with standard 19" 42U Rack with PDU
8.	Management	Dedicated 1GbE Management Port.
9.	Ports	2x USB 3.0 or more, 1x VGA
10.	Power Supply	Redundant and Hot Pluggable, 80 Plus Platinum or better certified power supply along with IEC 14 / 19 type Power cables
11.	Warranty	5 years onsite Comprehensive Warranty with NBD Support
12.	Remote Management	Remote management port, IPMI 2.0 or equivalent support with KVM and Media over LAN features with additional licenses if any. Should have support for Redfish API or equivalent for server management.

B) AI based Compute Nodes

Quantity		20 Nos.
Sr. No.	Specification	Description
1.	Processor	Dual x86_64 Processor CPU with minimum 32cores/64 threads Minimum base clock frequency of CPU 2.0 GHz., Support for Virtualization technology.
2.	Memory	Total 512GB Memory DDR 4, Min 3200 MHz or better with Advanced ECC with multi-bit error protection
3.	GPU	2 x NVIDIA A30 24GB CoWoS HBM2 PCIe 4.0-- Passive Cooling
4.	Hard Disk Drives	2 x 960GB SATA SSD or better
5.	HDD bays	8 x 3.5" or 2.5" hot-swap drive bay or more
6.	Expansion Slot	2 x PCIe Double-Width GPU Slots 4.0 x16, Min. 4 x PCIe Low-Profile Slots 4.0 x8 atleast
7.	Remote Management	Remote management port, IPMI 2.0 or equivalent support with KVM and Media over LAN features with additional licenses if any. Should have support for Redfish API or equivalent for server management.
8.	Network interface	At least 1 x 1 GbE Gigabit Ethernet LAN port, 2x 10G/1G Base-T Ethernet LAN ports
9.	InfiniBand	Single 100 Gbps InfiniBand HCA with hardware RDMA feature
10.	Ports	2x USB 3.0 or more, 1x VGA, 1x Management
11.	Graphics	Onboard Graphics
12.	Operating System	Support latest version of 64-bit RHEL/CentOS/Ubuntu/Open source Linux – latest stable releases

13.	Form Factor	4U or less, Rack mountable, Rail Kits to be included - compatible with standard 19" 42U Rack with PDU
14.	Power Supply	Redundant and Hot Pluggable, 80 Plus Platinum or better certified power supply along with IEC 14 / 19 type Power cables.
15.	Compliance	Server offered must comply with MEITY Notification: F. No.43/4/2019-IPHW-Meity (all revisions till date). Make in India certificate must be issued by Server OEM only. For Make in India Servers SMT of Motherboard/PCB must be done in India with Server OEM model/ part number embedded.
16.	Warranty	5 years onsite Comprehensive Warranty with NBD Support

C) AI Storage Nodes

Quantity		1 No.
Sr. No.	Specification	Description
1.	AINVMe Storage All Flash	Usable capacity of 160 TB based on NVMe SSDs (Data Center Grade or Enterprise class) with minimum 1DWPD endurance. The solution must be scalable to double the capacity for future requirements.
2.	Storage Platform	<ul style="list-style-type: none"> • The solution should deliver at least 99.99 uptime reliability. • The solution should provide a single namespace that can scale to EB's. • The solution shall support both Infiniband and Ethernet networking • The solution shall support Kubernetes Container Storage Integration (CSI) • The solution must be POSIX compliant and support NFS, SMB , S3 protocols and latest version of GPU Direct storage(GDS) • All system features and data services must be included in base license with no additional upcharges.
3.	Management	<ul style="list-style-type: none"> • The solution shall provide a full-featured GUI in addition to a command line interface. All tuning parameters must be accessible through the GUI • The management dashboard should have health and performance reporting, visualization, and overall system management functions are accessible using the command line interface (CLI) or the intuitive graphical user interface (GUI) management console for both hardware and filesystem. • -The solution shall allow for expansion of a filesystem without downtime • The solution shall support user authentication via LDAP and Active Directory.
4.	Performance	<ul style="list-style-type: none"> • All data and metadata is distributed evenly across the storage cluster delivering great performance, without any storage hot spots or I/O bottlenecks. Storage solution must deliver metadata performance of 100,000 files create/sec and 100,000 file stat/sec. • The solution shall provide a minimum bandwidth performance of at least 30 GB/second write and 40 GB/s read. • The solution shall provide a performance of 4K IOPS of at least 3 Million IOPs. It should be capable of holding metadata with a capacity to accommodate 1 Billion files. • The solution shall support both high transactional and bandwidth-based performance on a single filesystem. All performance specifications shall refer to a single filesystem that has not been further configured, adjusted, or tuned between runs of different benchmarks or I/O loads.
5.	Data protection	<ul style="list-style-type: none"> • The system shall support encryption at rest and in flight • There should be no single point of failure in the entire solution • The solution shall support the concurrent failure of up to 2 disks in a RAID group without loss of access to data and should have RAID 6 (8+2) disk configuration. • The solution shall support end-to-end data protection • The solution must use checksum to protect against bit flip or a drive fault to protect against silent data corruption • The solution must have intelligent rebuild mechanism to rebuild only the data that is actively stored on the failed server or SSD • All storage nodes/controllers must participate in the recovery process

		during rebuilds for faster recovery. The maximum permissible rebuilding time for disk(s) allowed is 12 Hours upon replacement.
6.	Benchmark:	<p>Open-source IOR or FIO benchmarks/Industry standard benchmarks:</p> <ul style="list-style-type: none"> ● Running on compute nodes with 1 MB transfer size (or in the multiples thereof) and file size double than total storage cache and I/O node memory. The total transfer size should be greater than the RAM available on the AI compute nodes. ● Benchmark must deliver the write throughput for Primary storage min. 30 GB/s ● Benchmark must deliver the read throughput for Primary storage min. 40 GB/s ● Benchmark must deliver 4K read performance of 3 million IOPS or higher <p>The solution shall support both high transactional and bandwidth- based performance on a single filesystem. All performance specifications shall refer to a single filesystem that has not been further configured. Bidder to submit the performance benchmarks as a part of the bid.</p>
7.	Warranty	5 years onsite Comprehensive Warranty with NBD Support

D) Storage with 2PiB usable capacity

Sr. No.	Specification	Description
1.	Architecture and Benchmark	Storage with 2PiB usable capacity with min. 7 GB/s Throughput (over Ethernet LAN interfaces only). Multiple clients can read/write from the same file simultaneously. Ability to read and write in parallel to same file or different files. Data striping across multiple I/O nodes. The solution should be capable of holding metadata with a capacity to accommodate 2 Billion files. Bidder to submit the performance benchmarks as a part of the bid.
2	Connectivity	Connectivity from NAS servers to compute nodes must be on HDR Infiniband and 10/25/40 Gbps Ethernet for Local LAN access. Connectivity between Storage servers to Storage enclosures with Redundant Connects & Links. Standard software features that creating/ managing volumes, storage & file system monitoring. If separate licensing is required for these features, it must be included in the proposed system.
3.	Balanced Architecture	Entire Solution must be in balanced with respect to number of NAS servers, There must be minimum one pair highly available NAS servers.
4.	High Availability	The solution should be highly available. Storage subsystem should be configured with no Single Point of Failure (SPOF). Power supply should be redundant and hot pluggable/replaceable, etc.
5.	Client License	Client licenses: Must support unlimited clients (compute/management servers) for accessing the file system.
6.	Warranty	5 years onsite Comprehensive Warranty with NBD Support

E) Login / Authentication Server

Quantity	1 No.	
Sr. No.	Specification	Description
1.	Processor	Dual x86_64 Processor CPU with minimum 24 cores/48 threads minimum base clock frequency of CPU 2.8 GHz.
2.	Memory	Total 512GB Memory
3.	Hard Disk Drives	5 x 12TB NLSAS 7.2K RPM 3.5" HDD & 2 x 960GB SATA SSD With hardware raid controller minimum supporting RAID 0,1,10 & 5 disk configurations
4.	HDD bays	12 x 3.5" or 2.5" hot-swap drive bay
5.	Expansion Slot	Min. 4 x PCIe Low-Profile Slots 4.0 x8 atleast

6.	Remote Management	Remote management port, IPMI 2.0 or equivalent support with KVM and Media over LAN features with additional licenses if any. Should have support for Redfish API or equivalent for server management.
7.	Network interface	1x 1 GbE Gigabit Ethernet LAN port, 2x 10G/1G Base-T Ethernet LAN ports 1 x 100G HDR HBA, 1 x Dual port 25G SFP+ or better
8.	Ports	2x USB 3.0 or more, 1x VGA, 1x Management
9.	Graphics	Onboard Graphics
10.	Operating System	Support latest version of 64-bit RHEL/CentOS/Ubuntu/Open source Linux – latest stable releases
11.	Warranty	5 years onsite Comprehensive Warranty with NBD Support
12.	Form Factor	2U Rack mountable, compatible with standard 19” 42U Rack with PDU. Server should be supplied with a compatible Railing Kit.
13.	Compliance	Server offered must comply with MEITY Notification: F. No.43/4/2019-IPHW-Meity (all revisions till date). Make in India certificate must be issued by Server OEM only. For Make in India Servers SMT of Motherboard/PCB must be done in India with Server OEM model/ part number embedded.

F) Network Switch for Desktop Connectivity

Quantity		12 Nos.
Sr. No.	Specification	Description
1.	Gigabit Smart Managed Switch	<ul style="list-style-type: none"> Two switches (Qty-02) with At least 40 x 100/1000 BASE-T ports & 8 mGig ports up to 10G (Total 48 - RJ45 ports). 2x 1/10/25G fixed uplinks over SFP28, Minimum (non-stack) standalone switching capacity 340 Gbps, forwarding rate 250 Mpps, 1U Rack mountable. Ten switches (Qty-10) with At least 48 x 100/1000 BASE-T ports (Total 48 - RJ45 ports). 4 x 1/10 G uplinks over SFP/SFP+ transceivers, Minimum (non-stack) standalone switching capacity 150 Gbps, forwarding rate 120 Mpps, 1U Rack mountable. PoE+ support with at least 740W of budgeted power for PDE, Dedicated Stacking port with 1 Mtr. stacking cable. Redundant power supply and power cables to be included. Advanced L2 switching and security features Supporting CDP, VTP, dual VLAN on access ports for data and voice, supporting DTLS tunneling for Cisco Wi-Fi Access Points Full feature support to cisco Access Points and Cisco IP Phones (over SCCP and SIP) Interconnect-able to existing Institutional Cisco LAN network over 1G/10G SFP+ or 10G/25G SFP28 Cisco transceiver modules.
2.	Warranty	1 year onsite Comprehensive Warranty with NBD Support
3.	Networking accessories	<ul style="list-style-type: none"> UTP patch cords (as required) for interconnecting the installed equipment's as the part of the supply. Additionally, 800 Factory molded, Cat 6 UTP Patch cords (2 Meter in length) is to be included. 2 Pairs of each (1 pair of Single mode and 1 of Multimode transceivers) for 1G/10G SFP+ & 10G/25G SFP28 Cisco transceiver modules: Total 4 Pairs (2 S.M. + 2 M.M.) equated to 8 transceiver modules.

G) Infini-band Switch HDR 200 for Primary Network

Quantity		1 No.
Sr. No.	Specification	Description
1.	InfiniBand 40-Port HDR Switch	HDR InfiniBand Switch, 40 QSFP56 ports, 2 Power Supplies (AC), x86 dual core, standard depth, P2C airflow, Rail Kit for rack mounting.

2.	Features	<ul style="list-style-type: none"> • 40 HDR 200Gb/s ports in a 1U switch • 80 HDR100 100Gb/s ports in a 1U switch • 16Tb/s aggregate switch throughput • Up to 15.8 billion messages per second • Enables offloading of collective operations from the CPU to the switch network • Quality-of-service enforcement • 1+1 redundant power supply • Makes the fabric capable with self-healing autonomy • Adaptive routing • InfiniBand router • Collective offloads • Virtual lane (VL) mapping (VL2VL)
3.	IB cables	Required number of Cables to connect all AI node servers, AI compute servers, Storages, NVMe storage with suitable length spanning across multiple Racks.
4.	Warranty	5 years onsite Comprehensive Warranty with NBD Support

H) Desktop Client Systems

Quantity		210 Nos.
Sr. No.	Specification	Description
1.	Processor	One Processor with minimum 6 Cores and minimum base clock frequency of 2.5 GHz for the latest generation of CPU from the OEM.
2.	Motherboard	Compatible with above processor
3.	RAM	8GB DDR4
4.	Storage	Minimum 1TB SATA HDD
5.	Graphic	Nvidia/AMD 2GB Graphic Card
6.	Chassis	Tower chassis
7.	Ports	Intel GbE LAN, USB 3.0/1/2, Audio Ports.
8.	Monitor	23.8" LED Monitor or higher
9.	Kb/Mouse	USB Wired Kb/mouse
10.	Warranty	5 years onsite Comprehensive Warranty with NBD Support
11.	Operating System	Centos 64 bit/ Ubuntu, latest and stable release available

I) Management Network Switch for AI

Quantity		1 No.
Sr. No.	Specification	Description
1.	Gigabit Switch	48 x 10/100/1000BASE-T ports
2.	Warranty	1 year onsite Comprehensive Warranty with NBD Support

J) Software

Sr. No.	Specification	Description
1	Job Scheduler/Queue management	<ul style="list-style-type: none"> • Slurm job scheduler Management supporting view information about Slurm nodes and partitions. • Should report resource usage for running or terminated jobs including individual tasks, which can be useful to detect load imbalance between the tasks. • Monitor and manage many jobs with identical resource requirements. • Supports the ability to define and schedule arbitrary Generic RESources (GRES). Additional built-in features are enabled for specific GRES types, including Graphics Processing Units (GPUs) and CUDA Multi-Process Service (MPS) devices, through an extensible plugin mechanism.

Other Mandatory Conditions: (Compliance to be given for each point)

1. Bidder needs to provide the OS/driver compatible to OS on All nodes.

2. Documentations: Detailed solution description along with architecture diagram, Product Data sheets, brochures etc. and detailed compliance to specification. The equipment / Product Data sheets, brochures should be publically available on OEM's Website.
3. Bidders should submit MAF from the all the OEMs of the products chose to bid.

Scope of Work:

1. Installation, Commissioning and integration of AI Based HPC Solution, network, and storage components in the supplied racks. Installation of client machines accessing the HPC, storage in the respective laboratories.
2. The bidder must implement/install/configure complete turnkey solution. Any additional hardware/component/sub-component/software required to complete the solution in its entirety and ensure production grade operations shall be the responsibility of the bidder.
3. The successful bidder is required to provide ONE On-site engineer for the facility management of the proposed deployment for the entire duration of warranty support (i.e. 5 Years) on 8 x 5 basis (8 Hours a day – 5 days a week) at no additional cost. The onsite engineer should essentially possess at least 2 years of experience in handling, troubleshooting and maintaining a similar computing infrastructure and should be at least B.E. / B.Tech. in Computers/ I.T./ Electronics Engineering.
4. Slurm Job scheduler Management supporting view information about Slurm nodes and partitions. Slurm should able to submit job in on containers directly.
5. Should report resource usage for running or terminated jobs including individual tasks, which can be useful to detect load imbalance between the tasks.
6. Solution should support Library and Compilers GCC, Compilation Environment Setup.
7. Installation of clustering software and management tool for cluster to be built upon.
8. Warranty: Five years comprehensive on Site Warranty on next business day support basis from the Bidder/OEM (except for Network Switches for Desktop Connectivity –which otherwise is required to be quoted with 1 year warranty). OEM and Bidder must not have been black listed or banned or debarred or stopped from doing business by any Govt. organization in the past. A notarized declaration must be submitted with the bid.
9. AI Frameworks Preloaded(CPU Optimised Tensor Flow, CPU Optimised Pytorch, CPU Optimised Theano, CPU Optimised Caffe, CPU Optimised Text2speech, CPU Optimised Mxnet , CPU Optimised CuDNN, Accelerator Optimised Tensor Flow, Accelerator Optimised Pytorch, Accelerator Optimised Theano, Accelerator Optimised Caffe, Accelerator Optimised Text2speech, Accelerator Optimised Mxnet, Accelerator Optimised CuDNN, Keras, MKL, DNN and support for a period of 5 years.
10. Containerization Support Containerizing Different Framework Building Customer Containers support for 5 years.
11. Inference Deployment & Performance Tools TensorRT.
12. Monitor and manage many jobs with identical resource requirements.
13. Supports the ability to define and schedule arbitrary Generic RESources (GRES). Additional built-in features are enabled for specific GRES types, including Graphics Processing Units (GPUs) and CUDA Multi-Process Service (MPS) devices, through an extensible plugin mechanism.
14. Bidder should provide details of power, cooling requirements, Rack Units (RU), number of racks and total rack space of the entire solution. Number of Racks required for the installation of the solution (excluding desktops) to be included as the part of supply such that after housing all equipment the rack loading should not be exceed 80% (including PDU sockets available in the Rack). This shall cater to closely knitted integration of the future augmented capacities.
15. Documentation and Manuals of all systems, configurations and policies implemented.
16. The storage related benchmarks, as elucidated earlier have to be submitted as a part of the technical bid and additionally demonstrated as a part of the technical acceptance after the installation and commissioning. The performance should be equal or better than those required in the tender, without which the entire solution shall not be considered as technically accepted.
17. Right to vary quantity: The institute reserves the right to vary the quantity of components/subcomponents in the above mentioned tendered requirement by 25% before placement of the final purchase order.