



INDIAN INSTITUTE OF TECHNOLOGY BOMBAY
MATERIALS MANAGEMENT DIVISION
Powai, Mumbai 400076

Ref. PR No.: 1000050994

Rfx. No.:6100002527

Item Description: Machine-Learning CPU/GPU Mixed Architecture Cluster

| Sr. No. | Technical Description of Item | | Technical Compliance (Yes/No) | Additional Information (if any) |
|-----------|---|---|-------------------------------|---------------------------------|
| 1. | Host / Management Node: Qty - 2Nos (in HA) | | | |
| 1. | Processors | Two numbers of 5th or Latest Generation Thirty-Two Core processor with x86 64 architecture each having: Minimum 2.1 GHz or higher Base Clock Frequency 5600 MT/s or higher L3 cache of 160MB or higher TDP should not be more than 270W Each Processor with 2100GF Rpeak Performance (with turbo frequency and hyper-threading feature disabled) | | |
| 2. | Motherboard | OEM motherboard with Minimum 32 nos. on-board DIMM slots | | |
| 3. | Memory | 256GB DDR5 Registered DIMMs with advanced ECC Scalable to 1024GB or higher without disturbing the existing memory modules | | |
| 4. | Storage Controller | 12G SAS RAID Hardware RAID controller with support for RAID 0/1/1+0/5 Proposed disk storage Controller should support Hardware root of trust and secure encryption and decryption of critical drive data/encryption via SED drives. | | |
| 5. | Hard disk | Minimum 2 x 1.92TB NVMe ENT SATA SSD in RAID 1 (OS Boot- mirrored) | | |
| 6. | Network Interface and Cluster HCA | 1 * Dual port 1G RJ45 1 * 1 GbE RJ45 Ports for Remote System Management Minimum 1 x Single-port NVIDIA ConnectX-7 400-Gigabit NDR400InfiniBand Adapter | | |
| 7. | Form factor | 1U only Rack Mountable form factor Chassis with sliding rail-kit | | |
| 8. | Industry standard Compliance | Energy Star ISO Certificate: for server quoted -9001, 14001, 27001 | | |

| | | | | |
|-----|---|---|--|--|
| | | EMI / EMC Certificate for quoted server model or substantially equivalent | | |
| | | BIS Certificate or substantially equivalent | | |
| | | CE/UL/FCC/IEC/EN certifications for quoted server model or substantially equivalent | | |
| 9. | Server Management | Zero Touch Provisioning (ZTP) using SSDP with remote access/ Unattended provisioning as per OEM | | |
| | | Pre-failure alert for key components like Processors, Memory, Internal storage etc | | |
| | | Should have dashboard for firmware baselines while performing minimum required firmware checks and highlighting out-of-compliance devices for updates with the selected firmware baseline | | |
| 10. | System Security | Immutable Silicon Root of Trust | | |
| | | Support for Commercial National Security Algorithms (CNSA) | | |
| | | Configurable for PCI DSS compliance | | |
| | | Chassis intrusion detection option | | |
| 11. | Firmware Security | For firmware security, system should support remote management chip creating a fingerprint in the silicon, preventing servers from booting up unless the firmware matches the fingerprint. This feature should be immutable | | |
| 12. | Embedded Remote Management and firmware security | Should support managing multiple servers as one via Group Virtual Media and Encrypted Virtual Media | | |
| | | Server should have security dashboard: displaying the status of important security features, the Overall security Status for the system, and the current configuration for the Security State and Server Configuration Lock features. | | |
| | | One-button/Commman Secure Erase designed to decommission/repurpose servers | | |
| | | Workload Performance Advisor- Provides server tuning recommendations to improve server performance being optional | | |
| 13. | Operating System Certified & Supported. | System quoted must be supported & certified for Redhat Linux operating system. This information should be listed on Redhat website under Redhat certified system catalog. | | |
| 14. | OEM | OEM of proposed Host Node must have minimum 10 entries as OEM in each list of India Top Supercomputer published from Jan 2022 to till the date of bid submission OR OEM of proposed Host Node must have minimum | | |

| | | | | |
|-----------|---|---|--|--|
| | | 10 entries as OEM in each list of worldwide TOP500 supercomputer published from Jan 2022 to till the date of bid submission. | | |
| 15. | System Fans | Hot plug and High-Performance redundant fans | | |
| 16. | Power supply | Hot plug redundant power supplies | | |
| 17. | Integration | Product Quoted must fully OEM factory integrated product | | |
| 18. | Warranty | 5 years of 24x7 onsite comprehensive warranty on all supplied hardware directly by the OEM. | | |
| 2. | Computing Nodes (Quantity 4) – High Memory (Large RAM) | | | |
| 1. | Processors | Two numbers of 5th or Latest Generation Thirty-Two Core processor with x86 64 architecture each having: | | |
| | | Minimum 2.1 GHz or higher Base Clock Frequency | | |
| | | 4800 MT/s or more | | |
| | | L3 cache of 160MB or higher | | |
| | | TDP should not be more than 270W | | |
| | | Each Processor with 2100GF Rpeak Performance (with turbo frequency and hyper-threading feature disabled) | | |
| 2. | Motherboard | OEM motherboard with Minimum 32 nos. on-board DIMM slots | | |
| 3. | Memory | 1024GB DDR5 Registered DIMMs with advanced ECC | | |
| | | Scalable to 2048GB or higher without disturbing the existing memory Modules | | |
| 4. | Storage Controller | 12G SAS RAID Hardware RAID controller with support for RAID 0, 1 | | |
| | | Proposed disk storage Controller should support Hardware root of trust and secure encryption and decryption of critical drive data/encryption via SED drives. | | |
| 5. | Hard disk | Minimum 2 x 1.92TB NVMe ENT SATA SSD in RAID 1 (OS Boot- mirrored) | | |
| 6. | Network Interface and Cluster HCA | 1 * Dual port 1G RJ45 | | |
| | | 1 * 1 GbE RJ45 Ports for Remote System Management | | |
| | | Minimum 1 x Single-port NVIDIA ConnectX-7 400-Gigabit NDR400InfiniBand Adapter | | |
| 7. | Form factor | 1U only Rack Mountable form factor Chassis with sliding rail-kit | | |
| 8. | Industry standard Compliance | Energy Star | | |
| | | ISO Certificate: for server quoted -9001, 14001, 27001 | | |

| | | | | |
|-----|---|---|--|--|
| | | EMI / EMC Certificate for quoted server model or substantially equivalent | | |
| | | BIS Certificate or substantially equivalent | | |
| | | CE/UL/FCC/IEC/EN certifications for quoted server model or substantially equivalent | | |
| 9. | Server Management | Zero Touch Provisioning (ZTP) using SSDP with remote access/ Unattended provisioning as per OEM | | |
| | | Pre-failure alert for key components like Processors, Memory, Internal storage etc | | |
| | | Should have dashboard for firmware baselines while performing minimum required firmware checks and highlighting out-of-compliance devices for updates with the selected firmware baseline | | |
| 10. | System Security | Immutable Silicon Root of Trust | | |
| | | Support for Commercial National Security Algorithms (CNSA) | | |
| | | Configurable for PCI DSS compliance | | |
| | | Chassis intrusion detection option | | |
| 11. | Firmware Security | For firmware security, system should support remote management chip creating a fingerprint in the silicon, preventing servers from booting up unless the firmware matches the fingerprint. This feature should be immutable | | |
| 12. | Embedded Remote Management and firmware security | Should support managing multiple servers as one via Group Virtual Media and Encrypted Virtual Media | | |
| | | Server should have security dashboard: displaying the status of important security features, the Overall security Status for the system, and the current configuration for the Security State and Server Configuration Lock features. | | |
| | | One-button/command Secure Erase designed to decommission/repurpose servers | | |
| | | Workload Performance Advisor- Provides server tuning recommendations to improve server performance being optional | | |
| 13. | Operating System Certified & Supported. | System quoted must be supported & certified for Redhat Linux operating system. This information should be listed on Redhat website under Redhat certified system catalog. | | |
| 14. | OEM | OEM of proposed Host Node must have minimum 10 entries as OEM in each list of India Top Supercomputer published from Jan 2022 to till the date of bid submission OR OEM of proposed Host Node must have minimum | | |

| | | | | |
|------------|---|---|--|--|
| | | 10 entries as OEM in each list of worldwide TOP500 supercomputer published from Jan 2022 to till the date of bid submission. | | |
| 15. | System Fans | Hot plug and High-Performance redundant fans | | |
| 16. | Power supply | Hot plug redundant power supplies | | |
| 17. | Integration | Product Quoted must fully OEM factory integrated product | | |
| 18. | Warranty | 5 years of 24x7 onsite comprehensive warranty on all supplied hardware directly by the OEM. | | |
| 3. | H200NVL system or substantially equivalent - Qty. - 02 No. | | | |
| 1. | Component | AI infrastructure and workloads from analytics to training to inference System - Specifications | | |
| 2. | Processors | Two numbers of 5th or Latest Generation atleast 56 core processorwith x86 64 architecture | | |
| 3. | NV Brdige | 900 Gb/s of bidirectional GPU-to-GPU bandwidth or 1.8TB/s bidirectional bandwidth between 4 x GPUs | | |
| 4. | GPU | 4Nos x GPUs that provide 564 GB total GPU memory | | |
| 5. | Performance | 13.36 peta FLOPS FP8 | | |
| 6. | Multi Instance GPU | Single GPU can be partitioned into as many as Up to 7 MIGs @ 16.5GB each GPU instances. | | |
| 7. | Security | GPU should support Confidential Computing | | |
| 8. | System Memory | Minimum 768 TB using 24 DIMMs system memory or Higher | | |
| 9. | CUDA Cores | Minimum 16000 or above, per GPU | | |
| 10. | Tensor Cores | Minimum 456 or above per GPU | | |
| 11. | Network | Network (storage and in-band management) card: 2 x NVIDIAConnectX-7 Dual Port NDR200 or substantially equivalent or Single port NDR IB/Ethernet Cards Each card provides the following speeds: Ethernet (default): 400GbE 200GbE, 100GbE, 50GbE, 40GbE, 25GbE, and 10GbE InfiniBand: Up to 400Gbps | | |
| 12. | OS Storage | 02Nos x 1.92 TB NVMe M.2 SSD in RAID 1 array | | |
| 13. | Internal Storage | 04Nos x 3.84 TB NVMe in RAID 0 array | | |
| 14. | Power requirements | Hot plug & redundant power supply 2+2 redundancy. 4Nos x 2400/2600 W @ 200-240V, 16A, 50-60 Hz. | | |

| | | | | |
|-----|---|---|--|--|
| | | C19/C20 Six (4) locking power cords that have been qualified for use with the system to ensure regulatory compliance | | |
| 15. | System Network (IPMI) | BMC (out-of-band system management) | | |
| | | 1 GbE RJ45 interface | | |
| | | Supports Redfish, IPMI, SNMP, KVM, and Web user interface | | |
| 16. | OS Support | Quoted OS (Ubuntu/RHEL/ROCKY server distribution with supporting packages.) should be under Enterprise support from OEM. | | |
| | | System Management – Provides active health monitoring and system alerts | | |
| | | Data Center GPU Management (DCGM) | | |
| | | This software enables node-wide administration of GPUs and can be used for cluster and data-center level management | | |
| | | The GPU driver | | |
| | | Docker Engine | | |
| | | Container Toolkit | | |
| | | Networking Open Fabrics Enterprise Distribution for Linux (MOFED) | | |
| | | Networking Software Tools (MST) | | |
| | | Cache file sd (daemon for managing cache data storage) | | |
| 17. | AI, HPC Software Containers and Required DL SDKs with Support | Some of the basic, SDK/library/containers we will use in the system: | | |
| | | CUDA toolkit, CUDA tuned Neural Network (cuDNN) Primitives | | |
| | | TensorRT Inference Engine | | |
| | | CUDA tuned BLAS (cuBLAS) | | |
| | | CUDA tuned Sparse Matrix Operations (cuSPARSE) | | |
| | | Multi-GPU Communications (NCCL) | | |
| | | Industry SDKs – NVIDIA Deep Stream, ISAAC, DRIVE, Nemo, Jarvis | | |
| 18. | Preinstalled AI frameworks | Installed optimized AI frameworks like Caffe, CNTK, Tensor flow, Theano, Torch with Docker containers for deploying Deep learning frameworks. | | |
| | | Pre-installed Deep learning GPU Training System for to train highly accurate deep neural network (DNNs) for image classification, segmentation, and object detection tasks. | | |
| 19. | | System should be scalable with multi node cluster. Software support & cluster tools to be | | |

| | | | | |
|-----------|---|--|--|--|
| | Scalability & Cluster software | supplied along with product. Full-stack reference designs with all of the leading Storage providers. | | |
| 20. | Warranty & Support | 5-year business-standard hardware and software support Warranty. | | |
| 21. | Cluster Management & Scheduler | Cluster management for node provisioning and monitoring needs to be included. | | |
| | | The Cluster Manager must support multiple of hardware vendors. The Cluster Manager must allow for the easy deployment and management of servers across multiple data centers, the public cloud, and edge locations as a single shared infrastructure through a single interface. | | |
| | | The Cluster Manager must provide the automatic repurposing of nodes between multiple HPC workload managers and Kubernetes clusters. | | |
| | | Highly Scalable Job Scheduling system to be included. | | |
| 4. | Computing Nodes (Quantity 18) – Regular Nodes. | | | |
| 1. | Processors | Two numbers of 5th or Latest Generation Thirty-Two Core processor with x86 64 architecture each having: | | |
| | | Minimum 2.1 GHz or higher Base Clock Frequency | | |
| | | 4800 MT/s or more | | |
| | | L3 cache of 160MB or higher | | |
| | | TDP should not be more than 270W | | |
| | | Each Processor with 2100GF Rpeak Performance (with turbo frequency and hyper-threading feature disabled) | | |
| 2. | Motherboard | OEM motherboard with Minimum 32 nos. on-board DIMM slots | | |
| 3. | Memory | 512GB DDR5 Registered DIMMs with advanced ECC | | |
| | | Scalable to 1024GB or higher without disturbing the existing memory Modules | | |
| 4. | Storage Controller | 12G SAS RAID Hardware RAID controller with support for RAID 0, 1 | | |
| | | Proposed disk storage Controller should support Hardware root of trust and secure encryption and decryption of critical drive data/encryption via SED drives. | | |
| 5. | Hard disk | Minimum 2 x 1.92TB NVMe ENT SATA SSD in RAID 1 (OS Boot- mirrored) | | |
| 6. | | 1 * Dual port 1G RJ45 | | |

| | | | | |
|------------|---|---|--|--|
| | Network Interface and Cluster HCA | 1 * 1 GbE RJ45 Ports for Remote System Management | | |
| | | Minimum 1 x Single-port NVIDIA ConnectX-7 400-Gigabit NDR400InfiniBand Adapter | | |
| 7. | Form factor | 1U only Rack Mountable form factor Chassis with sliding rail-kit | | |
| 8. | Industry standard Compliance | Energy Star | | |
| | | ISO Certificate: for server quoted -9001, 14001, 27001 or substantially equivalent | | |
| | | EMI / EMC Certificate for quoted server model or substantially equivalent | | |
| | | BIS Certificate or substantially equivalent | | |
| | | CE/UL/FCC/IEC/EN certifications for quoted server model or substantially equivalent | | |
| 9. | Server Management | Zero Touch Provisioning (ZTP) using SSDP with remote access/ Unattended provisioning as per OEM | | |
| | | Pre-failure alert for key components like Processors, Memory, Internal storage etc | | |
| | | Should have dashboard for firmware baselines while performing minimum required firmware checks and highlighting out-of-compliance devices for updates with the selected firmware baseline | | |
| 10. | System Security | Immutable Silicon Root of Trust | | |
| | | Support for Commercial National Security Algorithms (CNSA) | | |
| | | Configurable for PCI DSS compliance | | |
| | | Chassis intrusion detection option | | |
| 11. | Firmware Security | For firmware security, system should support remote management chip creating a fingerprint in the silicon, preventing servers from booting up unless the firmware matches the fingerprint. This feature should be immutable | | |
| 12. | Embedded Remote Management and firmware security | Should support managing multiple servers as one via Group Virtual Media and Encrypted Virtual Media | | |
| | | Server should have security dashboard: displaying the status of important security features, the Overall security Status for the system, and the current configuration for the Security State and Server Configuration Lock features. | | |
| | | One-button/command Secure Erase designed to decommission/repurpose servers | | |

| | | | | |
|-----------|--|--|--|--|
| | | Workload Performance Advisor- Provides server tuning recommendations to improve server performance being optional | | |
| 13. | Operating System Certified & Supported. | System quoted must be supported & certified for Redhat Linux operating system. This information should be listed on Redhat website under Redhat certified system catalog. | | |
| 14. | OEM | OEM of proposed Host Node must have minimum 10 entries as OEM in each list of India Top Supercomputer published from Jan 2022 to till the date of bid submission OR OEM of proposed Host Node must have minimum 10 entries as OEM in each list of worldwide TOP500 supercomputer published from Jan 2022 to till the date of bid submission. | | |
| 15. | System Fans | Hot plug and High-Performance redundant fans | | |
| 16. | Power supply | Hot plug redundant power supplies | | |
| 17. | Integration | Product Quoted must fully OEM factory integrated product | | |
| 18. | Warranty | 5 years of 24x7 onsite comprehensive warranty on all supplied hardware directly by the OEM. | | |
| 5. | Storage IO Nodes (Quantity 02 or higher as per solution demand) | | | |
| 1. | Processors | Two numbers of 5th or Latest Generation Sixteen Core processor with x86 64 architecture each having: | | |
| | | Minimum 2.8 GHz or higher Base Clock Frequency | | |
| | | 5200 MT/s or more | | |
| | | L3 cache of 37.5MB or higher | | |
| | | TDP should not be more than 195W | | |
| | Each Processor with 1400GF Rpeak Performance (with turbo frequency and hyper-threading feature disabled) | | | |
| 2. | Motherboard | OEM motherboard with Minimum 32 nos. on-board DIMM slots | | |
| 3. | Memory | 512 GB DDR5 Registered DIMMs with advanced ECC | | |
| | | Scalable to 1024GB or higher without disturbing the existing memory Modules | | |
| 4. | Storage Controller | 12G SAS RAID Hardware RAID controller with support for RAID 0, 1 | | |
| | | Proposed disk storage Controller should support Hardware root of trust and secure | | |

| | | | | |
|-----|---|---|--|--|
| | | encryption and decryption of critical drive data/encryption via SED drives. | | |
| 5. | Hard disk | Minimum 2 x 1.92TB NVMe ENT SATA SSD in RAID 1 (OS Boot- mirrored) | | |
| 6. | HBA | External Storage HBA with dual-port 12 Gbps SAS or 16 Gbps Fibre Channel, fully compatible with the proposed PFS storage solution. | | |
| 7. | Network Interface and Cluster HCA | 1 * Dual port 1G RJ45 | | |
| | | 1 * 1 GbE RJ45 Ports for Remote System Management | | |
| | | Minimum 1 x Single-port NVIDIA ConnectX-7 400-Gigabit NDR400InfiniBand Adapter | | |
| 8. | Form factor | 1U only Rack Mountable form factor Chassis with sliding rail-kit | | |
| 9. | Industry standard Compliance | Energy Star | | |
| | | ISO Certificate: for server quoted -9001, 14001, 27001 or substantially equivalent | | |
| | | EMI / EMC Certificate for quoted server model or substantially equivalent | | |
| | | BIS Certificate or substantially equivalent CE/UL/FCC/IEC/EN certifications for quoted server model or substantially equivalent | | |
| 10. | Server Management | Zero Touch Provisioning (ZTP) using SSDP with remote access/ Unattended provisioning as per OEM | | |
| | | Pre-failure alert for key components like Processors, Memory, Internal storage etc | | |
| | | Should have dashboard for firmware baselines while performing minimum required firmware checks and highlighting out-of-compliance devices for updates with the selected firmware baseline | | |
| 11. | System Security | Immutable Silicon Root of Trust | | |
| | | Support for Commercial National Security Algorithms (CNSA) | | |
| | | Configurable for PCI DSS compliance | | |
| | | Chassis intrusion detection option | | |
| 12. | Firmware Security | For firmware security, system should support remote management chip creating a fingerprint in the silicon, preventing servers from booting up unless the firmware matches the fingerprint. This feature should be immutable | | |
| 13. | Embedded Remote Management and firmware security | Should support managing multiple servers as one via Group Virtual Media and Encrypted Virtual Media | | |
| | | Server should have security dashboard: displaying the status of important security | | |

| | | | | |
|------------|--|--|--|--|
| | | features, the Overall security Status for the system, and the current configuration for the Security State and Server Configuration Lock features. | | |
| | | One-button/command Secure Erase designed to decommission/repurpose servers | | |
| | | Workload Performance Advisor- Provides server tuning recommendations to improve server performance being optional | | |
| 14. | Operating System Certified & Supported. | System quoted must be supported & certified for Redhat Linux operating system. This information should be listed on Redhat website under Redhat certified system catalog. | | |
| 15. | OEM | OEM of proposed Host Node must have minimum 10 entries as OEM in each list of India Top Supercomputer published from Jan 2022 to till the date of bid submission OR OEM of proposed Host Node must have minimum 10 entries as OEM in each list of worldwide TOP500 supercomputer published from Jan 2022 to till the date of bid submission. | | |
| 16. | System Fans | Hot plug and High-Performance redundant fans | | |
| 17. | Power supply | Hot plug redundant power supplies | | |
| 18. | Integration | Product Quoted must fully OEM factory integrated product | | |
| 19. | Warranty | 5 years of 24x7 onsite comprehensive warranty on all supplied hardware directly by the OEM. | | |
| 6. | PFS Storage | | | |
| 1. | Parallel File System | | | |
| 2. | Sever OEM/Bidder supported OEM supported GPFS or Lustre, Gluster FS or equivalent PFS with following specification:- 200 TB (usable in Hardware RAID 6 (8+2) configuration only) Parallel File System based storage. PFS should comprise of separate Metadata and Object storage servers (if required) with no single point of failure in HA to be configured. | | | |
| 3. | Metadata Storage: More than or equal to 5% of the Usable Storage space offered (using SSDs configured as RAID1 or Similar with one hot-spare or distributed metadata. | | | |
| 4. | Usable Storage(OST) : >= 500TB usable with Hardware RAID6 (8+2) Only (using 4TB or higher, 7.2K RPM NLSAS HDDs.) Software RAID is unacceptable | | | |
| 5. | Configured as RAID6 volumes with 2% hot-spare disks of the proposed storage. | | | |
| 6. | I/O Nodes - MDS nodes to be configured in Active-Passive and | | | |

| | | | |
|-----|---|--|--|
| | OSS Nodes to configured in Active-Active. | | |
| 7. | Fibre Channel external storage system with dual 12Gb SAS / 32/16Gb Fibre Channel active – active RAID controllers having a minimum of 48/64GB cache / Memory per array. | | |
| 8. | Offered storage shall have optional support for SSD / Flash based Cache for read operations. Offered Storage shall support at-least 8TB of Flash cache as an optional feature | | |
| 9. | Cache shall be backed up (Either through Battery, Capacitors or any other equivalent technologies) for an indefinite time or at-least 5 -7 years of time | | |
| 10. | Should support hot plug expansion and replacement of hard drives, redundant controllers, fans and power supplies | | |
| 11. | Offered storage shall support Raid 1, Raid 10, Raid 5 and Raid 6. Thin provisioning shall be supported on all the raid sets as well as with Flash cache | | |
| 12. | Offered storage subsystem shall support minimum 512 Logical units with LUN size support more than 120TB at storage controller level | | |
| 13. | The storage system shall support Enterprise SAS Drives, SSD and mid-line SAS drives of 7.2k rpm drives | | |
| 14. | The offered storage system shall also have support for all validating self-encrypting drives | | |
| 15. | Storage Array shall support at least 120/240 Enterprise SAS SFF drives or more than 100+ LFF drives | | |
| 16. | Offered storage shall be offered and configured with virtualization capability so that a given volume can be striped across all spindles of given drive type within a given disk pool. Disk pool shall support all listed raid sets of Raid 1, Raid 10, Raid 5 and Raid 6 | | |
| 17. | Raid processing shall be offloaded to a dedicated ASIC instead of CPU. In case vendor is not supporting it then vendor shall ensure that additional 32GB Cache is configured inside the storage arrays to offset the raid processing workload. | | |
| 18. | The offered storage shall also be licensed for Sub-Lun Data tiering in real time fashion across different types of drives within a given pool like SSD, SAS, Midline SAS etc. Licenses for same shall be offered for maximum supported capacity of the array. | | |
| 19. | Offered Storage shall have inbuilt performance management software and Configuration Dashboard shall show overall IOPS and MB/sec performance. | | |
| 20. | The offered Storage subsystem shall have controller-based Snapshot and Clone support. The storage subsystem shall be supplied with licenses for 512 Snapshots | | |
| 21. | The Offered Storage system shall support a minimum of 512 snapshots and 128 number of volume / clone copies as an optional feature. | | |
| 22. | The offered storage snapshot technology shall use Redirect on write technology. | | |

| | | | |
|-----------|---|---|--|
| 23. | Storage Array should have the capability of configuring hot swap disks and Global distributed Spare disks as well as dedicated spare disk to raid sets as an optional feature. | | |
| 24. | Offered storage shall support fan-in / Fan-out replication and it shall be supported to fan-out to at least 4 arrays. | | |
| 25. | The offered storage array shall have plug-in for VMware VCenter, Microsoft System centre as well as vStorage APIs (VAAI) for array integration as an optional feature. | | |
| 26. | Should support Windows server 2019 / 2022, Linux, VMware etc. | | |
| 27. | Offered storage shall have listed benchmark for performance of more than 700,000 IOPS for 100% random read and more than 14GB/sec for segmented sequential write operations. The vendor shall provide documentary proof for it. | | |
| 28. | The entire PFS solution must provide No Single Point of Failure (NSPOF) and capable of handling the loss of the following without interruption: -One Power Supply -One Fan - One SSD for MDT and Two HDDs for OST - One I/O ServerNodes The I/O server must have redundant paths to the storage. | | |
| 29. | All storage components (IO servers/Controller/Disk Enclosures) must be from the same server OEM so as to ensure seamless integration and homogeneity of the overall HPC solution. | | |
| 30. | 5 years of 24x7 onsite comprehensive warranty on all supplied hardware directly by the OEM. | | |
| 7. | High speed Primary Cluster Interconnect Network: | | |
| A- | NDR InfiniBand Switch | | |
| 1. | No. of Ports | 64 or more non-blocking InfiniBand 400 Gbps ports with at least 400 Gbps (one-way) bandwidth between every pair of ports. It should have at least 32 OSFP twin NDR InfiniBand ports | |
| 2. | Bandwidth | At least 51.2 Tbps of full bi-sectional switch bandwidth | |
| 3. | Latency | 100 ns or less | |
| 4. | Chassis | 1U Rack mountable chassis with rack mounting kit | |
| 5. | Power Supplies and Fans | Hot swappable redundant power supplies Hot swappable replaceable fan modules Air flow of fans should be from Power supply to Connectors side. | |
| 6. | Standards | IBTA specification 1.5 compliant | |
| 7. | Management | Switch should support Open Source OFED Subnet Manager for InfiniBand topology discovery and routing. | |
| 8. | Warranty | 5 years comprehensive warranty with all firmware and driver updates and bug fixes | |

| | | | |
|------------------------------|--|--|--|
| B- | NDR InfiniBand cables | | |
| 1. | NVIDIA active copper splitter cable, 800 (2x400) Gbps to 2x400Gbps, OSFP to 2xOSFP, 5m, fin to flat – 30Nos As per solution requirement. | | |
| 8. | Cluster Management Network: | | |
| 1. | All the compute nodes, master nodes and storage have to be connected 1Gig Ethernet switch for Cluster | | |
| 2. | Appropriate length cables to be provided for complete solution. | | |
| 3. | 5 Years onsite comprehensive warranty | | |
| 9. | Supported Cluster Manager and Job scheduler | | |
| Cluster Management Software: | Proposed cluster management tool to be fully deployed by its OEM/Bidder on-site. | | |
| | Proposed cluster management tool to be fully supported by its OEM/Bidder on-site. | | |
| | Proposed cluster management tool should have web based graphical remote access interface. | | |
| | Should be compatible to work on RHEL 8.0 or above/CentOS 6.x/ Fedora/Rocky Linux 8.0 or above/Ubuntu 22.04 or above | | |
| | Should be able to provision above operating systems to compute nodes. Should be able to auto-provision applications to compute nodes | | |
| | Should allow HPC management console to be accessible from any system in the network. | | |
| | Should support role-based access to the HPC system. | | |
| | Every role should be able to be remotely managed using the Graphic User Interface. | | |
| | Should provide profile-based and fully automated provisioning features. | | |
| | Proposed cluster management tool should have a graphic user interface. | | |
| | Should support Add/Modify/Delete compute nodes from GUI window. | | |
| | GUI (Web) based monitoring feature to be part of proposed solution. | | |
| | The licenses are required for entire existing compute nodes as well as the new ones to be procured through this tender. | | |
| | Should support major power management modules like bullpap, wti, apc_snmp, ether_wake, ipmilan, drac, ipmitool, ilo, rsa, lpar, bladecente as per the hardware proposed. | | |
| Job scheduler: | Bidder should propose a job scheduler which is fully supported by them. Job Scheduler proposed should only be OGE / Torque / Open PBS/SLURM or Licensed | | |

| | | | |
|------------------------------|--|--|--|
| | and supported version of PBS Pro. No other scheduler to be proposed. | | |
| Libraries, MPI and Compilers | OpenMP, MPI, C/C++ and FORTRAN compilers, GSL libraries should be installed. | | |
| | Bidder should implement Open MP and MPI network. | | |
| | Bidder should set up the Cluster for MPI Communication over IB/GigE | | |
| | Vendor specific InfiniBand stack on Linux OS, if available, should be supplied. | | |
| | Vendor specific MPI implementation on Linux OS should be supplied. | | |
| Scheduler | SLURM/Open PBS Scheduler configuration inputs have to be taken from the IIT-Bombay | | |
| Installation and Testing | Operating System – Latest stable version of Rocky Linux, Job scheduler installations as well as installation of common computational chemistry packages (example Gaussian-09/16, ORCA, Cp2k, VASP, MOLCAS, GAMESS, and others available at the department). Cluster functionality tests needs to be performed. | | |
| Hardware Installation | Successful vendors/OEM needs to install the cluster in the existing data centre and with the rack space provided. Since Racks space for the nodes are in place, no additional racks are solicited. | | |

General Eligibility Criteria

- 1.1 The bidder should be an OEM / Authorized Partner / Service Provider. A tender-specific with the Bid Number clearly mentioned on each Manufacturer Authorisation Letter (MAF) must be enclosed from Server/Storage/GUP server/Primary Interconnect. Failure to submit the same will result in disqualification.
- 1.2 The OEM/bidder should be ISO 9001 certified for Maintenance & System Integration (attach certificate).
- 1.3 Registration Number, LST/CST/WCT No., and PAN details must be provided.
- 1.4 The OEM must have had a direct presence in India for at least 10 years.
- 1.5 The OEM must have an average annual turnover of at least ₹20 Crores or equivalent over the last three consecutive financial years.
- 1.6 Bidders & OEM should not be suspended / blacklisted by any Central/State Govt Bodies including GeM and academic/research institutes in the last 7-years for any period in country. A declaration to be submitted in this regard.
- 1.7 Bidder must be capable of supplying the complete proposed hardware stack (Servers and Storage).
- 1.8 The quotation should specify the per-node price for each of the previously mentioned items, including compute and high-memory nodes, so that we can adjust the number of item/nodes as per the

| | | |
|--|--|--|
| project requirements/funds. There should not be any deviation from the price. | | |
| Technical Eligibility & Experience | | |
| <p>2.1 The OEM/Bidder should have supplied at least three HPC (CPU only) clusters in India, each with a minimum of 100 TF compute and 100TB PFS Storage. Submit list with contact details.</p> <p>2.2 Compute and storage OEM must have at least one 250TF HPC + 500TB PFS Storage installation reference in Government/Mod/Institutes organizations.</p> <p>2.3 Proposed CPU-only compute nodes must include LINPACK benchmark results by OEM (not CPU manufacturer). Logs and OEM undertaking required.</p> <p>2.4 The OEM of proposed must have: At least 10 entries in each of the India Top Supercomputers Published list (Jan 2022 onwards) OR At least 10 entries in each TOP500.org Published list (June 2022 onwards).</p> <p>2.5 The OEM must have at least 10 entries in each bi-annual Top500.org list over the last 3 years.</p> <p>2.9 The OEM must be well-established in the enterprise server and storage market for at least 10 years.</p> <p>3.0 The entire solution is expected to be delivered within 4 weeks' time once the PO is released.</p> | | |
| Product & Compliance Requirements | | |
| <p>3.1 For interoperability and seamless support, all servers/nodes/storage (Management, compute, High Memory Node, Normal CPU Node) must be from a single OEM.</p> <p>3.2 Product brochures/datasheets with full technical specs to be submitted. Link to OEM website for downloading brochures must be included.</p> <p>3.3 All hardware/software must be quoted with 5 years warranty (24x7x365 with 4-6 hr response) with a turn-key solution. OEM warranty undertaking to be enclosed.</p> <p>3.4 Warranty status must be available on OEM website via part/serial number.</p> <p>3.5 Power and cooling requirements with socket details to be submitted.</p> <p>3.6 Bill of Materials with OEM part numbers of all internal components must be submitted.</p> <p>3.7 Cluster should include installation & configuration of management tools, job schedulers (PBS Pro/Slurm), monitoring tools (Ganglia/IPMI), etc.</p> <p>3.8 Integration of CAE/CFD/NVH applications and MPI environment into the scheduler is mandatory.</p> <p>3.9 Data sheets must match RFP specs; any mismatch may lead to bid rejection.</p> <p>3.10 If requested by the evaluation committee, system demo (on-site) may be required.</p> <p>3.11 The vendor must ensure all proposed software (e.g., Gaussian, ORCA, VASP, etc.) run efficiently in a parallel cluster environment.</p> | | |

| | | |
|--|--|--|
| Service & Warranty Terms | | |
| <p>4.1 Bidder must provide a complete turnkey solution including installation and user training.</p> <p>4.2 The vendor must ensure 5 years of round-the-clock operation and maintenance from go-live date. Onsite warranty: 5 years.</p> <p>4.3 Installation must be done by certified and trained engineers.</p> <p>4.4 OEMs must submit an undertaking to support the bidder with updates and warranty fulfilment.</p> <p>4.5 Toll-free 24x7x365 customer care contact details of all OEMs must be submitted.</p> <p>4.6 OEM escalation matrix and contact details must be enclosed.</p> | | |
| Commercial Terms | | |
| <p>5.1 Bidder may quote in INR. Taxes and duties must be clearly mentioned.</p> <p>5.2 Provide per-node/component pricing. Submit commercial bid separately.</p> <p>5.3 Complete Scope of Work must be attached with the quotation.</p> <p>5.4 Payment Terms: Imported Components: Letter of Credit Indigenous Components: Post-supply and installation</p> <p>5.5 All required datasheets, brochures, compliance sheets, and supporting documentation must be submitted. Deviation without justification may lead to bid rejection.</p> | | |