INSTITUTE FOR PLASMA RESEARCH
(An Autonomous Institute of
Department of Atomic Energy, Government of India)
Near Indira Bridge; Bhat; Gandhinagar-382428; India

PART-I (B)
(TECHNICAL SPECIFICATIONS & COMPLIANCE SHEET)

Supply and installation of High Performance Computing System

IPR SCOPE for High Performance Computing System:

- Sufficient number of compute nodes must be quoted to achieve at least **24TF CPU** theoretical peak performances.
- Sufficient number of compute nodes with **GPU** must be quoted to achieve at least **6 TF GPU** theoretical peak Performance. (That is pure GPU and Not CPU+GPU).
- System/Compute Node should be able to run a single system Image (SSI) of64bit Red Hat Enterprise Linux OS for with 2 CPU-Processors with 28 Cores.
- Entire system must have comprehensive on-site warranty for 3 years. This is applied to all compute nodes, Compute nodes with GPU, Head/Master Node, all switches, Networking components, storage nodes and storage, Ethernet switches, IB switches etc.
- Appropriate number of commercial licenses with 3 years support/upgrade has to be quoted.
- OEM developed and owned Middleware/Management tools with latest version should be provided.
- Bidder should provide all necessary Cables, Connectors, optical drives, Components (Ethernet Cables, IB Cables, KVM cables etc.) required for quoted system.
- Bidder should provide Block Diagram of Quoted System.
- Technical compliance letter, detailed cluster diagram with datasheet of proposed solution to be submitted along with bid.

- **Note**: OEM refers to Original Equipment Manufacturer for **Master**/Head Nodes, Compute Nodes, compute Nodes with GPU, Cluster Management Software, Racks, storage server, storage.

Technical Specification of HPC Cluster

Hardware Specifications:

(1) Compute Node with CPU:

- **CPU**: Two Number of 14 core, 64 bit Intel Xeon E5-Series (2.3 GHz) or better processors, each core capable of executing 16 (sixteen) FLOPS or better per cycle .
- **Cache** : 35 MB or better
- **Memory**: 256GB DDR4 2133 MHz or better, optional **quote for 512GB DDR4 2133 MHz or better.**

- **Internal Disk**: 900 GB SAS disk with 10000 RPM or better
- **Form Factor**: Maximum 2 U height with rack/blade/dense form factor.
- **Infiniband**: Dual Port 4x FDR ports with 100% non-blocking architecture between nodes.
- **Network**: 2 X 1Gbps ports with PXE boot capability.
- **Power supply**: The solution should be configured with Hot plug Redundant Power supplies with efficiency.
- **Serviceability:** All the compute nodes should be individually serviceable without shutting down other compute nodes
- **Remote management Port:** At least one dedicated port for remote management.
- **GPU enabled**: All the compute nodes must be GPU enabled. Compute node must have at least 2 GPU slots with Gen 3.0 or better.

(2) Compute Node with CPU+GPU :

- **CPU**: Two Number of 14 core, 64 bit Intel Xeon E5-Series (2.3 GHz) or better processors, each core capable of executing 16 (sixteen) FLOPS or better per cycle.
- **Cache:** 35MB or better.
- **Memory**: 256GB DDR4 2133 MHz or better, optional **quote for 512GB DDR4 2133 MHz or better.**
- **Internal Disk**: 900 GB SAS disk with 10000 RPM or better
- **Form Factor**: Maximum 2 U height with rack/blade/dense form factor.
- **Infiniband**: Dual Port 4x FDR ports with 100% non-blocking architecture between nodes.
- **Network**: 2 X 1Gbps ports with PXE boot capability.
- **Power supply**: The solution should be configured with Hot plug Redundant Power supplies with efficiency.
- **Serviceability:** All the compute nodes should be individually serviceable without shutting down other compute nodes
- **Remote management Port:** At least one dedicated port for remote management.
- **GPU enabled**: All the compute nodes must be GPU enabled. Compute node must have at least 2 GPU slots with Gen 3.0.or better.
- **GPUs**: The compute node should be configured with Nvidia K40x GPU. Compute node has to be in proportion of GPU cards in 1:1 ratio of CPU: GPU. (Two GPU per node)

(3) Master /Head Node with High availability:

- **CPU**: Two Number of 14 core, 64 bit Intel Xeon E5-Series (2.3 GHz) or better processors, each core with capable of executing 16 (sixteen) FLOPS or better per clock cycle.
- **Cache:** 35 MB or better.
- **Memory**: 256GB DDR4 2133 MHz or better, optional **quote for 512GB DDR4 2133 MHz or better.**

- **Internal Disk**: 2 X 900GB SAS with 10000 RPM disks configured with RAID1. Software based RAID is not acceptable.
- **Form Factor**: Maximum 2 U height with rack/blade/dense form factor.
- **Infiniband**: Dual Port 4x FDR ports with 100% non-blocking architecture between nodes.
- **Network**: 4 X 1Gbps ports with PXE boot capability.
- **Power supply**: The solution should be configured with Hot plug Redundant Power supplies.
- **Serviceability:** All the compute nodes should be individually serviceable without shutting down other compute nodes
- **Remote management Port:** At least one dedicated port for remote management.
- **GPU enabled**: Nodes must be GPU enabled. Node must have at least 2 GPU slots.
- **GPUs**: One number of Nvidia K40x GPU in each node
- **Nodes:** Two nodes must be configured in pair of high availability to provide 100% redundancy.

(3) Networking:

- **Primary Interconnect – 4X FDR Infiniband :-**
    - 4x FDR - infiniband switch or better appropriate for HPC system with 100% non-blocking architecture. Redundant Power Supplies should be offered in all the switches.
    - Appropriate numbers of IB switches and IB cables, connectors etc. accessories must be quoted as per requirement will be supplied.

- **Admin and Console Network :**
    - GigE based Admin and Console network with appropriate managed L2 switches with redundant power supply to be supplied.
    - Appropriate number of CAT6 cables with suitable length to be supplied.

(5) Storage / Storage Nodes (Servers) with Parallel File System:

- Minimum 100 TB usable capacity with RAID 6 should be configured with parallel file system.
- All Storage nodes (Servers) required to implement Parallel File System should be in pair for High availability.
- All storage server should have dual port 4X FDR IB HBA.
- Storage should be connect to cluster's IB network via storage servers.
- The storage should be configured with at-least 1 GB/s WRITE and 1GB/s read performance (both READ and WRITE happening concurrently).
- Storage should be with dual controller with at least 8 GB total cache.
- Storage must be configured with redundant power supplies.
- Storage and Storage Server must be from same OEM and supported for 3 years.
- Offered solution must have no single point of failure.
- Solution must be fully compatible with quoted system.

Note: Optionally quote for 150 TB solution with above specification mentioned in no.(5)

(6) Management Console:

• 1U Rack mount LCD Keyboard and mouse with required cables and accessories.
• 8 port KVM SWITCH with cables compatible with the supplied system.

(7) Rack Enclosure:

- 42 U OEM rack of standard size with all required accessories to be supplied.
- All the rack components including power socket etc. must be fully compatible with quoted system.

Software Specifications:

(8) HPC Operating System and Software:

- **Operating System:** Latest version of 64 bit Red Hat enterprise Linux with 3 years Subscription and support. Licenses must be provided for all the nodes. (Master nodes, Compute nodes, GPU Compute nodes, Storage server nodes and other, if required.)
- **JOB scheduler**: Altair PBS pro suitable for all servers with 3 years support and subscriptions. It must support GPU jobs also.
- **Compilers**: Single user license of Intel cluster Studio with GPU support with 3 years support and subscription. Appropriate number of Licenses for all above software along with patches, upgrades etc. for 3 years.
- **Necessary tools for parallel programming on CPU and GPU like Open MP, Open ACC, Open CL and CUDA should be supplied.**
- **GPU libraries to be provided**
- **Supplied compliers and libraries must be suitable to offered solution.**

- **Cluster Management Software:**
  - Management software with appropriate licenses must be from same OEM.
  - Linux supported management software for monitoring and management of HPC cluster hardware like CPU, GPU, RAM etc.
  - Software should handle all the nodes provided in the solution.
  - It should support GUI/Web Based access.
  - It should provide proactive notifications alerts.

- **Parallel File System Software :**
  - Licensed Intel Sourced Lustre parallel file system.

Note: All software with latest version will be supplied. All software licenses must be perpetual.

(9) Delivery: Delivery and installation of the system at IPR-Bhat should be completed within 10 weeks from the date of purchase order/LOI.

(10) Implementation Schedule: Bidder will provide implementation schedule of the system.

(11) Documentation: Documentation on installation on commissioning of the system will be prepared and provided to IPR in soft and hard form.

(12) Training and Support:

- HPC –Cluster System Administration (Cluster Management, Configuration Monitoring, Job Scheduling etc.)
- Technical support for administration / maintenance (for both software and hardware level) for HPC.

(13) Application Porting / Migration:

- Bidder should provide technical help for migration of existing applications to the HPC.

(14) Warranty and Support:

- The 3 years warranty with onsite support will be effective from the date on which IPR accepts the system.

A) Mandatory documents to be uploaded by vendor during tendering:

a) Documents asked under "**PART-1 (A) Eligibility Criteria**"
b) Document with the **formula to achieve 24TF for CPU Nodes and 6 TF for pure GPU.**
c) Undertaking letter mentioning **"Make and model of Master/Head Nodes, Compute Nodes, compute Nodes with GPU, Cluster Management Software, Racks, storage server, and storage must be from same OEM"**
d) Implementation schedule.
e) Tender form/ Undertaking as per format available with the tender/enquiry

B) Acceptance Criteria :

Bidder must demonstrate theoretical peak performance as below:

- Theoretical Peak Performance along with method of arriving at the result.
- LINPACK Performance at least 70% for pure CPU nodes with Turbo OFF.
- LINPACK Performance at least 70% for CPU +GPU nodes with Turbo OFF.
- IOR /IO Zone benchmark for storage with at least 1 GB/s Write and 1 GB/s read. (Both READ and WRITE happening concurrently)