Request for proposals (RFP)

For the Supply of HPC system to the CSIR

RFP No. 3492/10/12/2021

Proposal Specification

Annexure A

1. HPC SYSTEM TECHNICAL REQUIREMENTS

The goal of this RFP is for the Bidders to provide a 4 PFlops HPC system for the CSIR as per the information below.

The solution must provide the following features and functionality:

a) At least 4 petaflops (PFlops) HPL rmax.
b) x86-64 CPU based system
c) Total solution must be maximum of 18 standard racks
d) Racks must be maximum height of 42U
e) Total solution must consume maximum of 1000kW
f) Solution must include all management nodes
g) 5 years support for the entire solution - software, and hardware
h) Monitoring and management tools
i) System management tools
j) At least 100Gb/s HPC interconnect per compute node with contention ratio of 1:1
k) NFS storage of 100TB and 100TB backup
l) Complete development environment with optimized compilers and libraries
m) Applications benchmark
n) PBSPro or Slurm resource manager
o) Cluster management software.
p) Integration of existing 3PB lustre storage – see Annexure M
q) Integrate the existing GPU nodes, Fat nodes, and NRE storage node – see Annexure M
r) The cooling system must integrate with existing water-cooling system
s) Bidder's experience in HPC implementation and number of HPC systems implemented

1.1. Management Nodes

The purposes of the management nodes are to enable the management of the other components in the system and to enable access to the hardware by the users. This will provide a login session for the users as a front-end. Users will submit batch scripts from the front-end as well as compile their applications, prepare data sets, and examine results.

All management nodes must have 1TB usable SSDs in a RAID10 configuration, with at least 24 cores and 2GiB memory per core and InfiniBand cards, unless specified.

The solution must provide:

a) 2× Login nodes with 10GBASE-SR Ethernet network cards.
b) 2× DTN nodes (Globus) with 100GBASE-SR4 Ethernet network cards.
c) 1× MedeA server with 10GBASE-SR Ethernet network card with a minimum of 64 cores and 4GiB memory per core. It must also have a minimum of 10TB usable SSDs in a RAID 10 configuration.
d) 1× Mslogin node with 10GBASE-SR Ethernet network card with a minimum of 64 cores and 4GiB memory per core. It must also have a minimum 10TB usable SSDs in a RAID 10 configuration.
e) 2× Cluster management (CM) nodes, supporting failover, with 10GBASE-SR Ethernet network cards and additional 1TB usable SSDs with RAID10 configuration. CM nodes will have usable 2TB SSD storage with RAID10 configuration.
f) 4× Visualisation nodes with 10GBASE-SR Ethernet network cards and a 16 GiB GDDR5 passively cooled Quadro RTX 5000 per node or another suitable card of equivalent or better level of specification. (See Annexure M)
g) 2× PBS/Slurm nodes, supporting failover, with 4GiB memory per core.

For the Management nodes, please provide details on;

a) Number and Type
b) Processor type and performance
c) Memory per node
d) Internal Disk Sub-System
e) Network connectivity

1.2. Compute Nodes

f) Compute nodes should be diskless
g) Compute nodes must consist of at least: 32 cores per node, 2 GiB of memory per core, and 8 memory channels per CPU socket.

For compute nodes, please provide details on:

a) Number and Type
b) Processor type and performance

c) Memory per node
d) Network connectivity

Model	
Processor	
Clock Speed (GHz)	
Number of instructions per clock	
Total number of compute nodes	
Number of processor sockets per node	
Number of processor cores per socket	
Number of memory channels per socket	
Number of threads per socket	
Memory size of nodes	
Cache and/or register sizes	
Other:	

(A copy of this table must be completed & returned with the tender)

1.3 Operating System

The Bidder shall propose a well-integrated and supported system software environment based on a RHEL 8 binary compatible operating system. The overall imperative is to provide users with a productive, high-performing, and reliable system software environment by which to use the system

The system shall include the stable supported release of Operating System (OS) on all visible partitions (nodes). Proprietary OS for management nodes can be proposed if necessary

1.4 Networking

The proposed system's high-speed interconnect network shall support high bandwidth, low latency, high throughput, and independent progress.

a) The proposed interconnect must have a total bandwidth per node of at least 100 Gb/s from one HCA.
b) The vendor shall describe the high-speed interconnect in detail, including any mechanisms for adapting to heavy loads or inoperable links.

Please provide details on the high-speed interconnectivity for your solution.

a) Network Topology;
b) Contention Ratio;
c) Bandwidth; and
d) Latency.

1.5 Storage and filesystem

1.5.1. Lustre Parallel FS

a) The CSIR has just completed the procurement of 3 PB Lustre file system that will be integrated as a parallel file system for this solution. The specifications for the 3 PB Lustre file system is available in Annexure M.
b) The bidder must propose in detail how they will integrate the 3 PB Lustre file system to their proposed solution
c) The bidder must propose all required equipment (hardware or/and software) for the 3 PB Lustre file system to be fully integrated to the proposed solution

1.5.2. NFS

a) NFS storage is required for users' home directories and storage of applications.
b) Bidders must provide two NFS servers, supporting failover, with at least 24 cores per node and 10GB memory per core.
c) Storage required for NFS is minimum 100TiB SSDs and a 100TiB disk based storage backup and recovery solution.

1.6 Integration of existing systems

Some of the CSIR legacy systems will need to be integrated to the proposed solution. Below is the list of systems that will need to be integrated: -

a) Lustre storage
b) NRE storage
c) GPU nodes
d) FAT nodes

1.7 System Management

1.7.1 System Cluster manager

The CSIR is using Bright Cluster Management (BCM) as their system management of choice. Bidders are however allowed to propose any system management software other than BCM. Should the bidder propose BCM software, extra licences must be added as required. Currently the CSIR is running Cluster Manager version 7.3 and CMDaemon 1.8 with a total of 1420 licenses of which 1415 are being used.

1.7.2 Workload manager

The CSIR is currently using PBSPro workload manager version 18.2. Bidders can only propose either PBSPro or SLURM.

1.8 System Monitoring and Operation

The system should provide a centralised capability for system administrators and operators to both remotely and locally monitor all major hardware and software components of the system and to detect and report normal and abnormal operation, resource utilization, and performance. Such monitoring capabilities should include all major software subsystems required for normal operation.

It is required that the proposed system has a web-based and CLI monitoring tool to monitor load and machine status

Please explain and outline, the monitoring capabilities of the system management tools with respect to:

a) Compute Nodes;
b) Network Infrastructure;
c) Storage Infrastructure.

1.9 Benchmarking

The CHPC has a large applications portfolio and will require benchmarking on a selected set of applications to be conducted on an HPC configuration representative of the proposed solution.

Bidders are required to provide benchmark data on all of these applications.

There are several goals associated with the benchmarking of applications, which are:

a) Compatibility;
b) Efficiency;
c) Performance; and
d) Scaling.

The applications in the benchmarking suite are currently in use at the CHPC and the data sets are of research interest.

Bidders should note that these applications would also form the acceptance tests for the facility once a preferred bidder has been identified. However, at this stage the CHPC seeks to obtain some information on the capabilities of the proposed solution with reference to the list of applications provided above and to the capabilities of the CHPC's current HPC solution.

Please provide the benchmarks for the above with any relevant system or application configuration information.

1.9.1 Benchmarking Summary

There are two types of benchmarks: synthetic and application. Both types of benchmarks are to be run on the proposed system.

1.9.1.1 Synthetic Benchmarks

A Complete HPC Challenge (HPCC) suite:

http://icl.cs.utk.edu/hpcc/

a) **HPL** – the **Linpack** TPP benchmark which measures the floating point rate of execution for solving a linear system of equations.
b) **HPCG –** the High Performance Conjugate Gradients benchmark is representative of a very wide range of practical HPC applications.
c) **DGEMM** – measures the floating point rate of execution of double precision real matrix-matrix multiplication.
d) **STREAM** – a simple synthetic benchmark program that measures sustainable memory bandwidth (in GB/s) and the corresponding computation rate for a simple vector kernel.
e) **PTRANS** (parallel matrix transpose) – exercises the communications where pairs of processors communicate with each other simultaneously. It is a useful test of the total communications capacity of the network.
f) **RandomAccess** – measures the rate of integer random updates of memory (GUPS).
g) **FFT** – measures the floating point rate of execution of double precision complex one-dimensional Discrete Fourier Transform (DFT).
h) **Communication bandwidth and latency** – a set of tests to measure latency and bandwidth of a number of simultaneous communications patterns; based on **b_eff** (effective bandwidth benchmark).

The HPCG Benchmark

hpcg-benchmark.org

1.9.1.2 Application Benchmarks (Input datasets to be supplied by the CHPC)

a) **OpenFOAM** (fluid dynamics).
b) **WRF** (weather).
c) **DL_POLY** (molecular dynamics)
d) **Quantum Espresso** (quantum mechanics)

1.9.2 Compatibility Benchmarks

Compatibility benchmarks are not intended to provide performance information, but to prove that these software products function on the proposed system. Ansys Fluent (fluid dynamics).

a) Siemens STAR-CCM+
b) BIOVIA Material Studio
c) Schrödinger

d) AMBER
e) Gaussian
f) VASP/MedeA (users use their own licenses at CHPC)

1.9.3 Results must be reproducible

Bidders must provide the CHPC with source modifications; build scripts, makefiles, run scripts, output files, important environmental settings, and node description(s). The benchmarks must be performed on hardware, which is similar to what will be delivered as part of the award. CHPC reserves the right to request that benchmarks be rerun before an award. CHPC is also interested in the discussions pertaining to the applications tuning process and any tools used to aid in that process to enable efficient parallelisation. Where bidders make use of tools, libraries, or drivers for the benchmark runs that are optimised or customised for their hardware, then those optimised software must be included as part of the proposal.

1.9.4 Not all applications need to run to completion.

Some of the applications in the test suite have a very long runtime with the data set of interest. The CHPC does not expect these applications to run to completion. The CHPC will indicate an acceptable termination point for a run.

1.9.5 Reporting

It is assumed that bibbers will run applications using various numbers of cores to show scalability. In addition to the files discussed above to ensure reproducibility, a summary of runs must be provided. For each run provide a one page description which contains:

a) Application Name;
b) Data Set Name;
c) Application mappings, for example:
    - Number of MPI tasks;
    - Number of nodes;
    - Number of cores; and
    - Thread settings.
d) Node information
e) Run time to designated termination point.
f) A URL link to the results (Dropbox or a similar service is acceptable).

2. IMPLEMENTATION

a) The installation and setup of the system must be performed at the CSIR facility in Rosebank, Cape Town.
b) The bidder must provide Project management plan and related information
c) The bidder must provide a support engineer on-site, available 8 am–4:30 pm Monday through Friday, with the exception of Public holidays, beginning on the project initiation date until system acceptance.
d) The support engineer will be expected to work closely with CSIR staff to develop, configure, and test the system until the acceptance stage and readiness to run the production workloads.
e) It is also expected that the knowledge of the system configuration, installation, and maintenance package will be transferred to CSIR staff during this period.

3. WARRANTY AND LICENCES

a) The cost of the proposed system shall include all warranties and licensing (if applicable) for all hardware and any software, delivered with the system. This must include hardware that is not manufactured by the primary Vendor. Warranties and licenses (if applicable) shall be for a period of 5 years of operation.
b) The start date of the year of full production operation is defined as the day of machine acceptance.

4. SKILLS TRANSFER AND TRAINING

a) Each proposal must include the implementation process that the Vendor will follow to work together with the CSIR staff.
b) Bidders must commit to on-site skills transfer for system administration of the equipment and software as part of their proposal.
c) Skills transfer shall commence before the implementation and prior to the beginning of the Acceptance Testing Period.
d) Training can be conducted virtually

The bidder must describe all proposed training (if not specified) and documentation relevant to the proposed solutions utilising the following methods:

a) Online documentation
b) Online training

5. ACCEPTANCE TEST SUITE

The Acceptance test suite will form part of the contractual agreement with the successful bidder and will be evaluated after the successful implementation of the proposed solution.

The Acceptance Test Suite will be comprised of the following items:

a) Versions of codes from the benchmark suite (from the above Synthetic and Application benchmark subsections) used to verify performance of the delivered system against proposed or reported performance as required in this document. The average of all benchmarks performed must be within 95% of the Bidder benchmarks and benchmark estimates provided in their proposal.

b) Integration of the 3PB Lustre will be verified using the IOR (both POSIX API and MPI-IO API) and mdtest benchmarks. Ref.: https://github.com/hpc/ior

c) System and workload monitoring utilities run by CSIR system staff.

d) Batch and interactive workload by “early users” from CHPC.

6.1. Acceptance Test Criteria

6.1.1. Duration

Once the system is installed and configured and the Bidder and the CSIR have agreed the system is ready, the system shall commence the Evaluation and Acceptance Test Period (ATP). The ATP is the first 14 consecutive calendar days, within the Evaluation Period, during which the system demonstrates compliance with the availability and reliability criteria below. If the system fails to meet one or more of the acceptance criteria after 30 calendar days from the date of Local Configuration Complete (also the commencement of the evaluation period), the CHPC may, at its option, terminate its Contract with the Bidder or extend the Evaluation Period.

6.1.2. Availability

The system shall achieve 96% system availability during the Acceptance Test Period. System availability is computed using the following terminology:

a) Total Time (TTot), time in minutes for the Acceptance Test Period—20,160 minutes.

b) Down Time (TDown), time in minutes when the system or a component thereof is inoperable and the system is unable to correctly run the Acceptance Test Suite due to a failure or malfunction in the Bidder supplied equipment or software, or because the system or component has been released to the Bidder for remedial services.

c) Down Time shall commence at the time when the CSIR contacts the Bidder's designated point of contact to report a failure. Down Time shall end when the system is returned to CSIR in an operable condition.

d) System Availability = Node total X (TTot – (non-functional nodes X TDown)/TTot

Down Time or system failure for reasons outside of the reasonable control of the Bidder, such as power outages, fires, floods, strikes, will not count against the Availability and Reliability requirement.

6.1.3. Reliability

a) There should be no failures that trigger downtime during the acceptance test. A failure is defined as any scheduled or unscheduled event that triggers Down Time, where TDown is when the system is inoperable and is unable to correctly run the Acceptance Test Suite.

b) Hardware support will require replacement of failed or faulty critical components by the next business day or sooner. A critical component is defined as one that is required to enable 90% of the system to be able to run projects successfully. Software support will require telephone and email response to problem calls and questions. Response to software support calls to be next business day or sooner.

c) Bidders should describe both the hardware and software support provided for the proposed system during the support period.

d) Bidders should specify the location of the nearest parts depot and explain the conditions under which support will be contracted out to a third party.

e) Bidders must provide information on how they anticipate delivering and supporting their proposed solution, paying particular attention to the personnel and skills available within their organisation.