

I.1: Commissioning of High Performance Computing Cluster, Kshitij-5 (क्षितिज-5)

High Performance Computing Cluster (HPCC) Kshitij-5 (क्षितिज-5) is the latest addition to the centralized scientific computing facilities available in RRCAT for scientific and engineering applications. Kshitij-5 combines the capabilities of Intel Xeon Skylake processors, NVIDIA's TESLA P100 Graphics Processing Unit (GPU) Accelerators, and Infiniband Enhanced Data Rate (EDR) 100 Gbps interconnect.

Kshitij-5, shown in the Figure I.1.1, is equipped with 64 compute nodes housed in 16 enclosures, eight rack mounted GPU nodes, eight rack mounted management servers, and 100 Gbps Mellanox InfiniBand (IB) switch as interconnect. Total 2048 Central Processing Unit (CPU) cores (from two Sixteen Core Intel Xeon Gold @2.60 GHz processors per compute node), along with 24 Tera Bytes (TB) aggregate memory, 16 GPUs (total 57,344 GPU cores) and a Storage Area Network (SAN) system with 300 TB raw capacity for storage are principle components of Kshitij-5.



Fig. I.1.1: Kshitij-5 HPCC.

Kshitij-5 is a blend of open-source software components and commercially available compilers and libraries. Various software components of this HPCC such as resource manager, parallel file system, user authentication service etc. are configured in fail-over mode.

The Lustre parallel file system has been configured to deliver high-throughput, high-availability and redundancy for enhanced Input Output (I/O) operations. This has been achieved with two nos. of Metadata Servers and four nos. of Object Storage Servers spread across four distinct nodes

Open-source Inter Process Communication Libraries such as

OPENMPI (openmpi-4.0), MVAPICH2 (mvapich2-2.3), and MPICH (mpich-3.3) are configured to support various types of parallel computations. Intel based compiler suite "Parallel Studio XE 2019" has been installed, which includes FORTRAN and C compilers, Math Kernel Library, Intel MPI etc. to support enhanced computations.

Web based scalable distributed monitoring system 'Ganglia' has been configured on Kshitij-5 and deployed on RRCATNet. Consolidated cluster usage in terms of Load, CPU, Memory, Network etc. are available through this monitoring tool. Figure I.1.2 shows typical screenshot of Ganglia monitoring system.

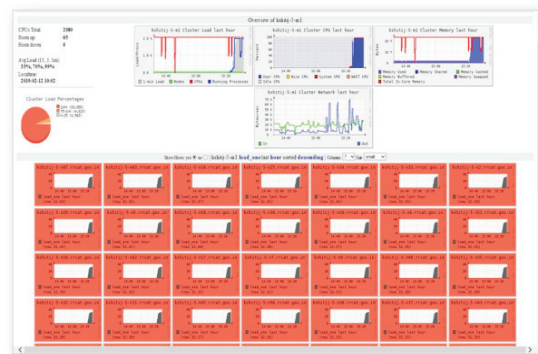


Fig. I.1.2: Load analysis of Kshitij-5 HPCC.

A) Benchmarking of Kshitij-5:

The sustained computing power delivered by Kshitij-5 HPCC was 190 Teraflops. This has been achieved via open-source parallel linpack and GPU enabled linpack benchmarks. The results obtained were further validated with commercially available Intel based parallel benchmarks.

B) Scientific applications ported to Kshitij-5:

As per requirement of users, many parallel applications have been successfully ported on Kshitij-5. The applications ported on Kshitij-5 include Pcrystal, MPPcrystal, VASP and WEIN2k for Human Resources Development Section, Parallel Elegant for Accelerator Physics Section, Epoch (epoch1d, epoch2d and epoch3d) for Laser Technology Division, and Orbit MPI for Accelerator & Beam Physics Section.

C) Centralized computing power of HPCCs in RRCAT:

After release of Kshitij-5, the aggregate centralized computing power of HPCCs available to scientists and engineers in RRCAT is 240 Teraflops.

Reported by:
Alpana Rajan (alpana@rrcat.gov.in)