

HPC Cluster - Kanad

Kanad, named after the great Hindu sage thought to have originated the idea of an atom, is an IBM iDataPlex cluster at IISER Bhopal. It is the institute's first High-Performance Computing facility serving the computational needs of several research groups on campus.

The iDataPlex racks are liquid-cooled and host 132 compute nodes (2112 processor cores) representing a theoretical peak performance of 44 Teraflops/sec. Each iDataPlex node runs on Intel E5-2670 series CPUs and has at least 64 GB memory. There are 4 high-memory nodes with 192 GB RAM for memory-intensive applications. The nodes are interconnected by 4X QDR InfiniBand network.

The cluster has access to 100 TB of storage divided in to a home (NFS) area of 60 TB and a parallel-file system of 40 TB as scratch. With a write through-put of over 2 GB/s this file system provides invaluable I/O speed when running applications.

Several software and libraries for scientific usage are installed on the cluster and have been optimized for performance.

Kanad runs the PBS job-scheduler and jobs are assigned in two major user pools :

Core user pool consists of users which heavy and regular need of computational resources. There are total of 127 nodes available for core users of which 4 are high-memory nodes.

Peripheral user pool consists of users with low resource requirements. There are a total of 5 nodes currently available for this pool. Resources are allotted following a core-based allocation policy in units of 8-cores and 32GB RAM.

IISER Bhopal students, researchers and faculty members can apply for accounts mentioning the user pool required along with a justification of the resources required. Application forms are available at [can be downloaded here](#).

Queues and Usage Policies :

The following queues are available for **Core** users :

Queue Name	Maximum Nodes	Maximum Walltime	Run Limit / Max. Nodes	Relative Priority
<i>interactive</i>	1	30 mins	4 jobs	1
<i>debug</i>	1 - 8	30 mins	4 jobs	1
<i>short1</i>	9 - 16	24 hours	90 nodes	2
<i>short2</i>	9 - 16	48 hours	64 nodes	3

HPC Cluster - Kanad

Queue Name	Maximum Nodes	Maximum Walltime	Run Limit / Max. Nodes	Relative Priority
<i>medium1</i>	4 - 8	48 hours	32 nodes	3
<i>medium2</i>	4 - 8	72 hours	64 nodes	4
<i>long1</i>	1-3	72 hours	15 nodes	4
<i>long2</i>	1 -4	120 hours	8 nodes	5
<i>hmq</i> (High Memory)	1 - 4	72 hours	4 jobs	3

The following queues are available for **Peripheral** users :

Queue Name	Range of cores	Maximum Walltime	Run Limit	Relative Priority
<i>p-queue</i>	8-16	48 hours	5 nodes	1

The *Run Limit* refers to the maximum total number of jobs (nodes) that can be running in a queue. *Queue Name* refers to the keyword to be used in a PBS job script to access the appropriate queue. Jobs will be accepted in the priority indicated by the last column above.

In addition the following policies need to be observed by users in the usage of the queues in order to ensure fair usage and avoid cancelation of their accounts :

1. Users will submit jobs only to through the queue. Running long production runs interactively on any node without using the queue is strictly forbidden.
2. Chaining of jobs on the interactive and debug queues are strictly forbidden.
3. User will not place jobs on hold for more that 4 hours. Such jobs shall be deleted by the systems administrator beyond the allowed limit.
4. A single core user can use for all jobs put together a total of **50 nodes**. For peripheral users this limit is currently set to 16 cores.
5. A single core (peripheral) user can have at most 3 jobs (2 jobs) waiting in the queue. Additional job requests will be automatically rejected by the scheduler.

The filesystems accessible to the users are divided into two main parts - Home (composing /home1 and /home2) and Scratch (/scratch). All areas are completely backed up on our tape-library (TSM-based).

Home Area :

Every user will be allocated a home area 1 TB for Core users and 500 GB for Peripheral Users, by default. Additional space can be allocated based on adequate justification. The

HPC Cluster - Kanad

location of this area will be *e.g. /home1/userid* and can also be accessed by the environmental variable \$HOME.

Scratch Area :

Every user will be assigned their own directory in the scratch area - *e.g. /scratch/<userid>*. This is also accessible by the environmental variable \$SCRATCH. Currently, there is no quota limit for any user in the scratch area. However, please be advised that files in the scratch area more than **14 days** old shall be purged. So users are advised to backup dated files.

For more information on software and hardware available on the cluster kindly contact the HPC Users Committee.