

General information

The main use case of the VLBI HPC cluster is doing VLBI correlations. However some of the cluster resources can be used for other computationally intense tasks e.g. simulations. For an overview of the cluster capabilities see Cluster Layout below.

If you want to have access to the cluster please get in touch with H.Rottmann (see Section [Contact](#))

Contact

Who	What	Phone
Helge Rottmann	cluster administrator	123
Jan Wagner	cluster administrator	365
Rolf Märtens	operator	220

Using the cluster

- Users need to get an account first (see H. Rottmann / J.Wagner).
- Once your account has been set-up you can log-into the cluster via ssh to `frontend`.
- `/home` contains your home directory and is visible on all cluster nodes
- `/cluster` contains software installations and is visible on all cluster nodes.
- each compute node has a local storage of 1TB available under `/scratch`
- Printing to any CUPS printer of the MPIfR should be possible from `frontend`.
- If you need to install software please consult with the administrators first.

Rules of conduct

- **Do not store dataset in your home directory.** If you need storage capacity talk to the cluster administrators first.
- Use only those cluster nodes that have been assigned to you by the cluster administrators.
- When done with your project tell the administrators so that the nodes can be used for other tasks.
- VLBI correlation always has first priority. When asked by the administrators users have to interrupt their jobs and return the assigned nodes back to the correlator (this happens rarely)

Preinstalled Software

The cluster has a number of software suites pre-installed, e.g. AIPS, ParselTongue etc. If you need additional software that might be of general interest talk to the administrators to have it installed on the cluster.

Special instructions exist for:

- [AIPS](#)
- [ParselTongue](#)
- [difmap](#)
- [ehtim](#) (EHT imaging library)

Activating software version

Some of the software, e.g. openMPI exists on the cluster in various versions. The available versions can be queried by

```
modules avail
```

Sample output:

```

backintime/1.0.40          knem/1.1.2          openmpi/1.10.0/gcc/
4.8.2 openmpi/1.10.1/gcc/5.2.0 openmpi/3.0.0/gcc/4.8.2
gcc/5.2.0                 munge/0.5.11      openmpi/1.10.1/gcc/
4.8.2 openmpi/2.1.2/gcc/4.8.2 slurm/15.08.3

```

To activate a certain version, e.g.:

```
module load openmpi/1.10.1/gcc/4.8.2
```

Cluster Specs

- The cluster has 68 compute nodes, each equipped with 20 cores. These are named `node01` to `node68`. The nodes can be accessed from `frontend`.
- All nodes are equipped with 64 GB of memory
- Each node has a scratch disk (~1TB)
- The cluster interconnect is realized with FDR Infininband @56 Gbps.
- User data storage is typically available under `/data11/users` (consult with the admins if you need storage capacity)
- Software available:
 - gcc 4.2.1 (g++, gfortran)
 - gcc 3.3.5 (g++, gfortran; under `/cluster/gcc/3.3.5/bin/`)
 - OpenMPI (various versions)
 - Intel Performance Library (IPP)
 - The SLURM batch system could be configured and enabled if need arises

Cluster Layout



io11, io12, io16	Infiniband Switch	4x Infiniband Switch, eportal1	4x Infiniband Switch, meta02
vbackup2	io03, io04, fxmanager2	nodes 29 - 48	Nodes 1 - 28
IO10 / 46 TB	nodes 49-52, meta01, io13-15	fxmaanger, appliance	frontend
IO05 - IO09 / 40 TB	Nodes 53 - 68	4 CASA nodes	IO01, IO02