

The Department of Computer Science hosts a computing cluster for scientific workloads. Please also consider the [HPC Services offered by our ZID](#) as they may suit your requirements better.

Hardware

The System is located in the IFI server room (3W06) in the ICT Building. It consists of 17 Nodes with the following specification:

```
Motherboard: ASRock Fatal1ty X399 Professional Gaming
CPU: AMD Ryzen Threadripper 2950X WOF
Memory: G.Skill D4128GB 2400-15 Flare X K8
SSD: WD Black SN750 500 GB, NVMe 1.3 , Read: 3470 MB/s, Write: 2600 MB/s
```

Nodes **gc1-gc8** offer

```
GPU: 4x ZOTAC GeForce RTX 2070 Blower, 8 GB (GDDR6, 256 Bit)
```

Nodes **gc9-gc16** offer

```
GPU: 4x ASUS GeForce RTX 2070 SUPER TURBO EVO, 8 GB (GDDR6, 256 Bit)
```

Node **gc17** offers

```
GPU: 1x NVIDIA TITAN RTX, 24 GB (GDDR6, 384 Bit)
```

The headnode (1x ASUS GeForce RTX 2070 SUPER TURBO EVO, 8 GB) has 2TB additional storage capacity that is reachable from all nodes.

Network

The nodes and storage are interconnected through a 10 Gigabit Ethernet Switch. The upstream connection of the headnode is Gigabit Ethernet.

System Configuration

- All nodes run the latest version of Ubuntu Server LTS which is 18.04.2 (Bionic Beaver).
- All nodes have the same home directories mounted.
- SLURM is used as job scheduler.
- [IFI Auth](#) is used as authentication backend.

Authorization & Authentication

To request access to ifi-cluster, users have to subscribe to the [mailing list](#).

After successful registration you can log into the system with your IFI credentials using SSH on **ifi-cluster.uibk.ac.at**.

Storage

Your **home directory** /home/name.surname should not exceed 100GB. Permission should be 700 in case you do not share it.

```
$ chmod 700 /home/name.surname
```

A 18TB HDD space is available to the GPU Cluster as a **scratch space**. It uses the 1Gbit network.

Mount point is : /scratch on each node

For better usage, please create a directory for yourself using the same username you have for the cluster and set accessibility only to you.

(in case you need to share a directory, give 750 or 755 permissions for group/others)

```
$ cd /scratch
$ mkdir name.surname
$ chmod -R 700 /scratch/name.surname
```

The **IFI-NAS storage** is mounted at /ifi-NAS/[your_group]. Information about IFI-NAS is available [here](#).

Mounting an external resource:

First create your own directory:

```
mkdir /mnt/users-mount/[your_dir]
```

Then mount the external resource:

```
sudo mount server:/resource /mnt/users-mount/[your_dir]
```

When finished, please do not forget to unmount it and delete the directory:

```
sudo umount /mnt/users-mount/[your_dir]
rm -rf /mnt/users-mount/[your_dir]
```

Installed Software Packages

Ubuntu Distribution Packages:

- Python 2.7
- Python 3.6
- gcc 7.4
- GNU Make 4.1

Loadable Modules:

- Open MPI 4.0.0 (/software-shared)
- mpich 3.3.1 (/software-shared)

Modules

[Environment Modules](#) can be used to modify a users environment. Use this to dynamically load modules in your sbatch script.

Basic commands:

```
module avail -- list available modules
module show [module name] -- show details about one module
module help [module name] -- show help from one module
```

Example usage in sbatch script:

```
module load openmpi
module unload openmpi
```

Topology

To be linked here soon...

Usage

Show running Jobs

queue

Submit Job

Batch Job

Write a configuration file similar to this basic example:

[cluster-test.sh](#)

```
#!/bin/bash -l
#SBATCH --partition=IFIgpu
#SBATCH --job-name=firstDemo
#SBATCH --mail-type=BEGIN,END,FAIL ##optional
#SBATCH --mail-user=max.mustermann@uibk.ac.at ##optional
#SBATCH --account=your_group ##change to your group
#SBATCH --uid=your_username ##change to your username
#SBATCH --nodes=2
#SBATCH --ntasks-per-node=1
#SBATCH --mem=4G
#SBATCH --time=0-00:30:00
#SBATCH --output slurm.%N.%j.out # STDOUT
#SBATCH --error slurm.%N.%j.err # STDERR
srun /bin/hostname | /usr/bin/sort
```

Option	Description
--partition=	specifies on which partition you want to run your job. Available partitions are: IFIall - for CPU computation, IFIgpu - for GPU computation , IFItitan - for large GPU computation on a NVIDIA Titan GPU Card
--job-name=	you can define a name for your job (this field is optional)
--mail-type=	to receive an e-mail if desired. Valid type values are NONE, BEGIN, END, FAIL, REQUEUE, ALL, +others see official documentation (this field is optional)
--mail-user=	specify the users e-mail (it could be any UIBK e-mail). if --mail-type is specified to other then NONE and --mail-user is missing e-mails will be send to local user on the IFI cluster (check it with "mail" command on headnode). For system wide mail forwarding please check here
--account=	specify the group you are belonging to
--uid=	specify your_username
--nodes=	define on how many nodes your job should run. The scheduler will allocate the number of nodes to your job.
--ntasks-per-node=	Request that ntasks be invoked on each node
--mem=	Specify the real memory required per node - default units are megabytes. Different units can be specified using the suffix [K M G T]
--time=	Set a limit on the total run time of the job allocation. If the requested time limit exceeds the partition's time limit, the job will be left in a PENDING state (possibly indefinitely).

Option	Description
<code>--output=<filename pattern></code>	The batch script's standard output will be directly to the file name specified in the "filename pattern". By default both standard output and standard error are directed to the same file.
<code>--error=<filename pattern></code>	The batch script's standard error will be directed to the file name specified in the "filename pattern".

The command that you want to run:

```
srun /bin/hostname | /usr/bin/sort
```

Run the file with **sbatch --test-only <<filename>>** to check the syntax and with **sbatch <<filename>>** to submit it.

Other valid and useful options:

Option	Description
<code>--exclusive[=user]</code>	The job allocation can not share nodes with other running jobs (or just other users with the "=user" option).
<code>--gres=<list></code>	Specifies a comma delimited list of generic consumable resources (GPUs). The format of each entry on the list is "name:type:count" (i.e. <code>--gres=gpu:4</code>)

There is also an

example script

for C-Compilation.

Use the official website documentation for details: <https://slurm.schedmd.com/sbatch.html>

Interactive Job

```
srun /bin/hostname
```

Interactive shell session

```
srun --nodes=1 --nodelist=gc3 --ntasks-per-node=1 --time=01:00:00 --pty bash -i
```

will get you a shell on gc3 node for one hour

Cancel Job

```
scancel <<jobid>>
```

The basic commands can be found in the [SLURM User Guide](#)

Live Monitoring Tool

The live monitoring tool for the cluster can be found [HERE](#)

Contact

Contact ifi-sysadmin@informatik.uibk.ac.at if you have problems or further questions. For administration purposes there is also the internal documentation.

IFI Internal Chat

IFI Internal Chat Room for the GPU Cluster can be accessed at [Matrix Chat](#).

From:

<https://ifi-wiki.uibk.ac.at/> - **IFI Wiki**

Permanent link:

<https://ifi-wiki.uibk.ac.at/public/ifi-cluster?rev=1619638775>

Last update: **2021/04/28 21:39**