

# S3DF

S3DF is SLAC Shared Scientific Data Facility. [Official Web-Page](#).

This scientific computing cluster contains shared computing resources of SLAC and can be used by our group members.

## Log-in

You have to get a user account. Connect via ssh.

```
ssh username@s3dflogin.slac.stanford.edu
```

The log-in nodes have very limited functionality. You have to continue to an interactive node.  
For example

```
ssh iana
```

## Theory Group Specific Resources

At the moment our resources on S3DF are still limited as we have not purchased any yet.  
This will change in fall 2023. Each user is currently equipped with a home directory of 28 GB (default).

We have a group directory located at "/sdf/group/epptheory/" with 10TB of available storage.  
There we install common software packages to be used by the Theory Group members.  
If you want additional software to be install contact Alex or Bernhard.

Currently, the following tools are installed:

```
gcc (v12 including g++, gfortran), Cuba, CLN, GINAC, LHAPDF, mysql
```

To make use of the installed software automatically you might want to add to your .bashrc file in your home directory the following lines.

```
export PATH=/sdf/group/epptheory/bin:/sdf/group/epptheory/Programs/gcc-12/bin:$PATH
export LD_LIBRARY_PATH=/sdf/group/epptheory/lib:/sdf/group/epptheory/Programs/gcc-12/lib64:/sdf/group/epptheory/Programs/mysql/lib:$LD_LIBRARY_PATH
export LD_RUN_PATH=/sdf/group/epptheory/Programs/gcc-12/lib:$LD_RUN_PATH
```

## SLURM

To use S3DF two options are in principle available.

First, there are interactive nodes like "iana" that allow you to log in and just run programs from the command line.  
The theory group has purchased two interactive nodes, which should arrive in fall 2023.

## Submitting Jobs

Second, the batch farm can be used to submit jobs via SLURM.

Here, we are benefitting from shared resources that are owned by others but are currently idle.  
We can submit our jobs to the farm but we have very low priority to run given we have no purchased resources ourselves.  
Nevertheless, this works quite well.

To submit a job you can either use command line. For example to run the script "test.out", do the following:

```
sbatch --ntasks 1 --cpus-per-task 4 --mem-per-cpu 4g -t 24:00:00 -o output.txt -e error.txt test.out
```

The above will allocate 4 CPUs for a time of 24 hours with 4 GB of memory per CPU to run your script and write the output and errors into the output.txt and error.txt files.

The above options are optional and just listed as an example.

To submit a batch script you can create first a file (e.g. "submit.slurm").

```
#!/bin/bash
#SBATCH --mem-per-cpu=4g
#SBATCH --cpus-per-task=4
#SBATCH --ntasks=1
#SBATCH -t 24:00:00
#SBATCH -o output.txt
#SBATCH -e error.txt
test.out
```

Next you can submit this script using the command

```
sbatch submit.slurm
```

## Managing jobs

`sacct`

See running jobs.

`scancel #id`

Cancel your job with job-id #id.

More [here](#).