

Maxwell for Photon Science users

This page is for external Photon Science User with a scientific account. Photon Science members please read the page [Photon Science on Maxwell](#).

The Interactive Photon Science resource in Maxwell

The Maxwell cluster contains two distinct Photon Science areas: an entirely interactive part and a conventional batch part (see below). To use the interactive part, connect to [desy-ps-cpu.desy.de](#) or - if you need a GPU for your calculations - [desy-ps-gpu.desy.de](#). The nodes are **not** accessible from outside the DESY network! For remote access to these nodes you first need to connect to [desy-ps-ext.desy.de](#) and ssh from there to [desy-ps-cpu.desy.de](#) or create an ssh tunnel on your local host.

Summary

- If you cannot access the system (not a member of netgroup @hasy-users), please contact fs-ec@desy.de or the beamline scientist connected to your experiment.
- If you're missing software or for support, contact maxwell.service@desy.de or fs-ec@desy.de or talk to your beamline scientist for advice.
- dCache instances `petra3` and `flash1` are mounted on `max-fsc/g (/pnfs/desy.de/ ; no Infiniband)`
- A `gpfs-scratch` folder is available by `/gpfs/petra3/scratch/` (approx 11T, no snapshot, no backup, automatic clean-up [will come]). It's scratch space in a classical sense, i.e. for temporary data. It's shared among all users and thus you're ask to check and to clean up / free scratch disk space if no longer needed.

[desy-ps-cpu.desy.de](#) is a very limited shared resource. You can only use a few cores and limited memory. If you use more than appropriate you will affects lots of colleagues. Some applications are very effective in crashing a node. So it might well be that a machine goes turn terminating all processes of all users. It happens much more frequent on the interactive nodes than on batch nodes. And if it happens on batch nodes there will be a single user affected any noone else. Batch nodes are hence much more efficient to use than interactive nodes: you get a full machine (or multiple machine) exclusively for your job. You can use all cores and memory available. For any serious data analysis or simulations the batch component of the Maxwell cluster is your best choice!

The Photon Science Batch resource for external users

As a first step login to [desy-ps-cpu.desy.de](#) and check which Maxwell resources are available for your account using the `my-partitions` command:

my-partitions

```
[@max-fsc ~]$ my-partitions

      Partition  Access  Allowed
groups
-----
-----
      all        yes    all          # <---- this one will be yes if any of the
resources below are available
      cfel       no     cfel-wgs-users
      cms        no     max-cms-uhh-users,max-cms-desy-users
      cms-desy   no     max-cms-desy-users
      cms-uhh    no     max-cms-uhh-users
      cssb       no     max-cssb-users
      epyc-eval  no     all
      exfel      no     exfel-wgs-users
      exfel-spb  no     exfel-theory-users,school-users
      exfel-th   no     exfel-theory-users
      exfel-theory no     exfel-theory-users
      exfel-wp72 no     exfel-theory-users
      fspetra    no     max-fspetra-users
      grid       no     max-grid-users
      jhub       no     all
      maxwell    no     maxwell-users,school-users
      p06        no     max-p06-users
      petra4     no     p4_sim
      ps         no     max-ps2-users
      psx        yes    max-psx2-users # <---- look for this one!
      uke        no     max-uke-users
      upex       no     upex-users,school-users
      xfel-guest no     max-xfel-guest-users,p4_sim
      xfel-op    no     max-xfel-op-users
      xfel-sim   no     max-xfel-sim-users

* As member of netgroup hasy-users you can (also) use max-fsc and max-fsg for interactive logins
* As member of netgroup psx-users you can (also) use desy-ps-cpu and desy-ps-gpu for interactive logins
```

If it says "yes" for partition "psx" you are ready to go. If so you will also see a "yes" at least for partition "all". If not: get in touch with FS-EC! Let's assume that you've got the psx-resource. The psx-resource also offers great additional opportunities for [remote login](#). It entitles you to use the max-display nodes to connect to the maxwell cluster via your [browser](#) or the fastx-client.

Apart from that: If you have an application, which is started by a script called my-application, and doesn't require a GUI, you can simply submit the script as a batch-job:

sbatch

```
[@desy-ps ~]$ sbatch --partition=psx --time=12:00:00 --nodes=1 my-application
Submitted batch job 1613895

# the job might already be running
[@desy-ps ~]$ squeue -u $USER
      JOBID PARTITION  NAME  USER ST  TIME  NODES NODELIST(REASON)
      1614464      psx  my-app user  R   0:06      1 max-wn052
# Status of the job
                                     R: running. PD: pending
```

This works for any application smart enough not to strictly require an X-environment, matlab, comsol, ansys, mathematica, idl and many others can be executed as batch jobs. To make it more convenient you can add the SLURM directives directly into the script:

SBATCH script

```
[@desy-ps ~]$ cat my-application
#!/bin/bash
#SBATCH --partition=psx
#SBATCH --time=1-12:00:00      # request 1 day and 12 hours
#SBATCH --mail-type=END,FAIL  # send mail when the job has finished or failed
#SBATCH --nodes=1             # number of nodes
#SBATCH --output=%x-%N-%j.out # per default slurm writes output to slurm-<jobid>.out. There are a number
of options to customize the job
[...] # the actual script.
```

The email-notification will be sent to <user-id>@mail.desy.de. That should always work, so you don't actually need to specify an email-address. If you do, please make sure it's a valid address. For further examples and instructions please read [Running Jobs on Maxwell](#).

If you think that it's much to complicated to write job-scripts or if you can't afford to invest the time to look into it: we are happy to assist. Please drop a message to maxwell.service@desy.de, we'll try our best.

Running interactive batch jobs

If you absolutely need an interactive environment, X-windows features like a GUI, there are options to do that in the batch environment. For example:

SALLOC

```
# request one node for 8 hours:
[@desy-ps ~]$ salloc --nodes=1 --time=08:00:00 --partition=all
salloc: Pending job allocation 1618422
salloc: job 1618422 queued and waiting for resources
salloc: job 1618422 has been allocated resources
salloc: Granted job allocation 1618422
salloc: Waiting for resource configuration
salloc: Nodes max-p3ag022 are ready for job

# now you got a node allocated. So you can ssh into the node
[@desy-ps ~]$ ssh max-p3ag022
[@max-p3ag022 ~]$ # run your application!
[@max-p3ag022 ~]$ exit # this terminates the ssh session, it does NOT terminate the allocation
logout
Connection to max-p3ag022 closed.
[@desy-ps ~]$ exit
exit
salloc: Relinquishing job allocation 1618422
# now your allocation is finished. If in doubt use squeue -u $USER or svview to check for running sessions!
```

There are a few things to consider:

- Interactive jobs with salloc easily get forgotten, leaving precious resources idle. We do accounting and monitoring! If the node utilization is low on the ps partition over some time you will start to get annoying emails. If it stays like this, your allocation might be removed. On the all partition we will let your job run nevertheless. If there is an urgent need the job will be terminated by competing jobs. So to be nice to you colleagues please considering using the all partition for interactive jobs. Keep the time short: there is hardly a good reason to run an interactive jobs for more than the working hours. Use a batch job instead.
- Terminate allocations as soon as the job is done!

Hardware Environment

For an overview of available compute nodes have a look at:

- [Hardware in the psx partition](#)
- [Hardware in the maxwell partition](#)
- [Hardware in the all partition](#)

General Environment

Personal environment

The home-directory is on a network-storage (currently hosted on GPFS) located in /home/<user>:

```
[user@desy-ps ~]$ echo $HOME
/home/user
```

This working directory resides on GPFS and thus you'll have the same "home" on all nodes of the Maxwell cluster.

Your home-directory is subject to a non-extendable, hard quota of 20GB. To check the quota:

```
[user@max-p3a001 ~]$ mmlsquota max-home
      Block Limits
Filesystem type      KB      quota      limit  in_doubt  grace |      File Limits
in_doubt  grace Remarks  files  quota      limit
max-home  USR      76576      0  10485760      0    none |      520      0      0
0         none core.desy.de
```

Software environment

The systems are set up with CentOS 7.x (RedHat derivate) and provide common software packages as well as some packages in particular of interest for photon science (e.g ccp4 [xds], conuss, fdmnes, phenix or xrt).

An overview of software available can be found here (no confluence account necessary for reading): <https://confluence.desy.de/display/IS/Software>

Some of the software can be executed directly, some packages are NOT available out-of-the-box and have to be loaded via module before you can use them. Which package to load (if necessary), how to load it and additional comments are given in confluence as well (see for example subsection <https://confluence.desy.de/display/IS/xds>). In case software is missing please let us know (email to maxwell.service@desy.de).

Storage environment

The GPFS core is mounted on the system as

```
/asap3/petra/gpfs
```

Underneath you'll find data taken recently at PETRA-Beamlines, sorted in subfolders by beamline with additional substructure of year/type/tag, e.g. data in 2015 at beamline p00 with beamtime AppID 12345678 would reside in `/asap3/petra/gpfs/p00/2015/12345678/`, followed by the directory tree as given/created during beamtime. Folder "logic" is same as during beamtime, i.e. assuming you've been named participant for a beamtime and are granted access to the data (controlled by ACIs [access control list]), you can read data from subfolder "raw" and store analysis results in `scratch_cc` (temporary/testing results) or processed (final results). For further details concerning folders and meaning please see subsections in confluence.desy.de, topic ASAP3.

The folder exported read-only to the Beamline PCs foreseen for documentation, macros etc. can be found in folder `/asap3/petra/gpfs/common/<beamline>`.

If you need large(r) amount of TEMPORARY space, please get in touch with maxwell.service@desy.de. There is beegfs space available on the cluster which is well suited for temporary data.