

Maxwell for CFEL (DESY)

The CFEL resource in Maxwell are own by a subgroup of CFEL-DESY and usage of the resources (compute & storage) is usually restricted to that group. If in doubt get in touch with the CFEL DESY admins.

The CFEL resource in Maxwell consists of a few interactive login nodes, some batch resources in the cfel partition and substantial GPFS storage. The compute nodes vary about the CPU-type and availability of GPUs. Please have a look

at [Maxwell Hardware](#) page, and the [limits](#) and [constraints](#) applying.

Interactive Login Nodes

To login in to CFEL part of the maxwell cluster you have different:

- **ssh max-display.desy.de**: will connect you to one of the display nodes. FastX might your better choice. Please have a look at the [Remote Login](#) and the [FastX documentation](#).
- **ssh max-cfel.desy.de**: will connect you to one of the interactive login nodes. It's a load-balanced alias. Use theses nodes to compile, develop or test applications.
- **ssh max-cfelg.desy.de**: will connect you to one of the interactive login nodes with GPUs. The nodes are equipped with dual K20X GPUs.
- **ssh max-wgs**: will connect you to the generic login node.
- Unless you use hardware-specific compiler-flags, compilation and job-submission can done from any of the login nodes.
- **Please note**: max-display.desy.de is directly accessible from outside. All other login nodes can only be reached by first connecting to bastion.desy.de.

Login nodes are always shared resources sometimes used by a large number of concurrent users. Don't run compute or memory intense jobs on the login nodes, use a batch job instead!

The CFEL Batch resource in Maxwell

As a first step login to one of the login nodes and check which Maxwell resources are available for your account using the my-partitions command:

```
my-partitions

[ @max-exfl ~ ]$ my-partitions

      Partition  Access  Allowed
groups
-----
----
resources below is "yes!"
cfel          yes    cfel-wgs-users    <----- look for this one as an CFEL
member
      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      yes    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    yes    maxwell-users,school-users    <----- might be granted if you have
suitable applications
      p06        no     max-p06-users
      petra4     no     p4_sim
      ps         no     max-ps2-users
      psx        no     max-psx2-users
      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
```

If it says "yes" for partition "cfel" you are ready to go. If so you will also see a "yes" at least for partition "all". If not: get in touch with CFEL DESY admins! Let's assume that you've got the cfel-resource. The cfel-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

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

# the job might already be running
[@max-cfel ~]$ squeue -u $USER
          JOBID PARTITION    NAME    USER ST       TIME  NODES NODELIST(REASON)
          1614464      cfel   my-app   user  R         0:06     1 max-cfel004
# 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

```
[@max-cfel ~]$ cat my-application
#!/bin/bash
#SBATCH --partition=cfel
#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:
[@max-cfel ~]$ 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-cfel005 are ready for job

# now you got a node allocated. So you can ssh into the node
[@max-cfel ~]$ ssh max-cfel005
[@max-cfel005 ~]$ # run your application!
[@max-cfel005 ~]$ exit # this terminates the ssh session, it does NOT terminate the allocation
logout
Connection to max-cfel005 closed.
[@max-exfl ~]$ 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!
- 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!

Other Maxwell Resources

Being member of CFEL and maybe having access to the cfel partition doesn't need to be the end of the story. If you have parallelized applications suitable for the Maxwell cluster you can apply for the Maxwell resource like everyone else on campus. Please send a message to [maxwell.service@desy.de](mailto:service@desy.de) briefly explaining your use case. Being also a user of the European XFEL you might also have access to the UPEX partition. You can easily distribute your job over the partitions:

multiple partitions

```
[@max-fsc ~]$ cat my-application
#!/bin/bash
#SBATCH --partition=cfel,maxwell,upex,all
#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 partition will be selected from cfel OR maxwell OR upex OR all starting with the highest priority partition. So your job will run on the cfel partition if nodes are available, on the maxwell partition if nodes are available and finally on the all partition if none of the other partitions specified have free nodes. Keep in mind that you should however select the partition according to the type of work you are doing. And a job can never combine nodes from different partitions, so check the [limits applying to partitions](#).

To check availability of nodes and characteristics use `sinfo` (<https://slurm.schedmd.com/sinfo.html>)

sinfo

```
[@max-display001 ~]$ sinfo -p cfel -o "%10P %.6D %8c %8L %12l %8m %30f %N"
PARTITION  NODES CPUS   DEFAULTT TIMELIMIT  MEMORY  AVAIL_FEATURES  NODELIST
cfel       6 64    1:00:00  14-00:00:00  512000  INTEL,V3,E5-2698,512G  max-cfel[003-008]
cfel       6 32    1:00:00  14-00:00:00  256000  INTEL,V3,E5-2640,256G  max-cfel[011-016]
cfel       4 40    1:00:00  14-00:00:00  256000  INTEL,V4,E5-2640,256G  max-cfel[017-020]
cfel       4 32    1:00:00  14-00:00:00  256000  INTEL,V2,E5-2650,256G,GPU,K20X max-cfelg[003-006]

[@max-display001 ~]$ sinfo -p cfel -o "%10P %.6D %10s"
PARTITION  NODES JOB_SIZE
cfel       20 1-6
```