

# Using the GWGD Scientific Compute Cluster

by Azat Khuziyakhmetov and Marcus Boden

Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

Burckhardtweg 4, 37077 Göttingen

Fon: +49 551 39-30000 [gwdg@gwdg.de](mailto:gwdg@gwdg.de) [www.gwdg.de](http://www.gwdg.de)

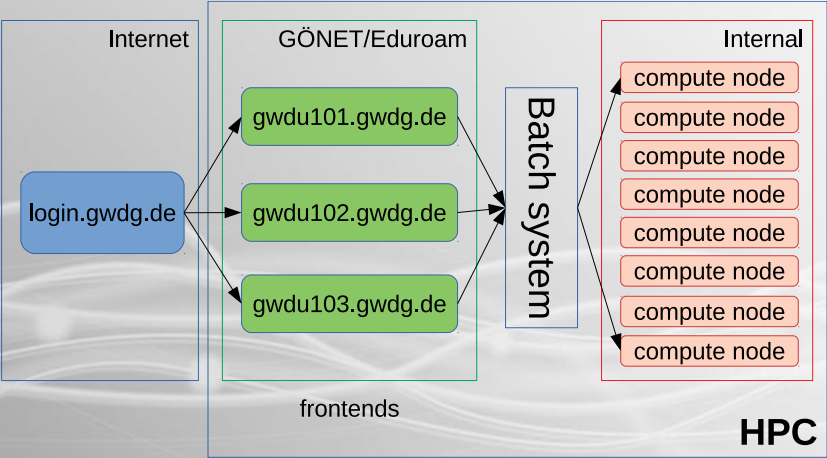
# Outline



- 1 Connecting to the frontends
- 2 The most important Linux commands
- 3 Specifics of GWDG HPC cluster
- 4 Preparing the environment with “modules”
- 5 Compiling Software
- 6 Using Slurm - Basics
- 7 Using Slurm - Advanced
- 8 Getting Help

## Section 1

### Connecting to the frontends



- gwdu101 and gwdu102 (new name: login-mdc.hpc.gwdg.de):
  - Cascade Lake Intel Silver 4214
  - processor features identical to agqXXX, agtXXX, ampXXX
  - new nodes in gpu and medium partitions
  - access to /scratch1 (formerly known as /scratch)
- gwdu103 (new name: login-fas.hpc.gwdg.de):
  - Broadwell Intel E5-2650 v4
  - processor features identical to dfaXXX, dmpXXX, dgeXXX, dteXXX
  - nodes in fat, medium and gpu partitions
  - access to /scratch2

From the Internet connect to “login.gwdg.de” first in similar way as shown below. Afterwards to the frontend node.

## You need SSH keys to connect to the cluster

- Linux or OS X:  
`ssh gwdu101.gwdg.de -I {GWWDG-USERID} -i {YOUR-KEY}`
- Windows: in newer versions you can use native “ssh” in power shell or download *putty.exe* from <https://www.putty.org>
  - ➔ Run it. Enter “gwdu101.gwdg.de” in *hostname*
  - ➔ In the menu SSH->Connection->Auth select your private key and click open
  - ➔ Select “Yes” to trust the connection
  - ➔ Login as: {GWWDG-USERID}

```
The authenticity of host 'gwdu101.gwdg.de (134.76.8.101)' can't...  
ECDSA key fingerprint is SHA256:sIJNEepmILeEq/7Zqq4HCtpTM8L98ar...  
Are you sure you want to continue connecting (yes/no/[fingerprint])?
```

## Section 2

### The most important Linux commands

`ls` list the current directory you are in

`cd` change directory



- List the current directory you are in, “ls”
  - List the “hidden” files (beginning with “.”) too, “ls -a”
  - All files in an extended manner, “ls -la” or just type “l”
- Let’s look at three lines of the output

```
drwxrwxrwx  3 akhuziy users    4096  4. Apr 17:29 test
-rw-r--r--  1 akhuziy users     283 24. Sep 2019 Info.txt
lrwxrwxrwx  1 root root         23  Jul 22 12:10 passwd -> /etc/passwd
```

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

ten permission flags:

- 1 directory flag, “d”: directory, “-”: normal file, “l”: symlink
- 2,3,4 read, write, execute permission for **U**ser (Owner of the file)
- 5,6,7 read, write, execute permission for **G**roup
- 8,9,10 read, write, execute permission for **O**thers

# Changing the language, what if I don't understand German



```
> echo $LANG
de_DE.UTF-8
> rm test
rm: reguläre leere Datei "test" entfernen?
> export LANG=en_US.UTF-8
> rm test
rm: remove regular empty file 'test'?
```

For persistent English language, put it in your ".profile":

```
echo 'export LANG=en_US.UTF-8' >> ~/.profile
```

**touch** create file / update timestamps

**other file ops** cp, rm, mv, mkdir, rmdir, ln

**htop** display Linux processes

**ps** display current processes, imp. opt. a [all sessions], u [owner], x [all], w [wide], ww [even wider]

**du** display file space usage, du -hs

**du** display file size summary, du -hs -apparent-size

**df** display filesystem usage, df -h, df -hl

- Files attributes (mode bits) can be changed with `chmod`
- `chmod` can be used in two ways:
  - ➔ user friendly form:  
u (user) g (group) o (others) a (all)  
`chmod a+r {file}`, `chmod g=rwx,o+r {file}`
  - ➔ tell the mode bits:  
`chmod 744 {file}`

- 0-7 are 3 bits: 111 → 7
- same order, like in dir listing: r,w,x

000 0 → --- no read write or execute allowed

001 1 → --x (last bit is set)

010 2 → -w- (middle bit is set)

011 3 → -wx (last 2 bits are set)

100 4 → r-- (first bit is set)

101 5 → r-x (first and last bits are set)

110 6 → rw- (first and second bits are set)

111 7 → rwx (all 3 bits are set)

chmod 456: owner - read; group - read and execute; others - read and write

- In sum we have 9 bits now in 3 groups (user, group, others)
- But there is a 4th group: SUID/SGID/sticky-bits
- SUID/SGID means that the called program will run with the UID or GID of the owner
  - ➔ e.g. if the program owns root and has SUID set, you run the program as root
  - ➔ `chmod u+s {file}`, or `chmod g+s {file}`,  
`chmod a+s {file}` would set both
  - ➔ Since we are normal users on the system, this is very seldom needed.

- SGID and sticky-bit for directories is more relevant for you, if you open a directory for colleagues to write (`chmod g=rwx {dir}`)
  - ➔ the SGID (`chmod g+s {dir}`) changes the group of a newly created file in this directory to the group of the directory, so all members of the group have the same permissions on all files.
  - ➔ the stick-bit prevents others from deleting files, they do not own. (`chmod +t {dir}`)
  - ➔ e.g. if you create a file, others cannot delete it, even though they have write permission to the directory.

- **nano**, vi/vim, mcedit, joe

For most commands you can read the manual pages, just type “man {COMMAND}”.

The prompt is also called “Shell” with its own commands and functions. We are using the `bash` shell. Type “man bash” to get an impression about the power and flexibility of that shell.



- **nano**, vi/vim, mcedit, joe

For most commands you can read the manual pages, just type “man {COMMAND}”.

The prompt is a so called “Shell” with built-in commands and functions. We are using the “bash”. Type “man bash” to get an impression about the power and flexibility of that shell.

Where the system gets all the commands we learned today?

Bash searches all paths in the environment variable **PATH**.

```
gwdu101:84 15:03:22 ~ > echo -e ${PATH//:/:\n}
/opt/slurm/bin:
/usr/lib64/qt-3.3/bin:
/usr/local/bin:
/usr/bin:
/usr/local/sbin:
/usr/sbin
```

For our first Shell script we need additional information

- “`mkdir -p /tmp/${USER}/XXXXXXXX`” will create a unique directory, e.g. `/tmp/akhuziy/XymeK4nq` and echo it to stdout
- To store an output of a program in a variable, we write “`TEMPDIR=$(mkdir -p /tmp/${USER}/XXXXXXXX)`”

Let's write a little Shell script

For our first Shell script we need additional information

- “`mkdir -p /tmp/${USER}/XXXXXXXX`” will create a unique directory, e.g. `/tmp/akhuziy/XymeK4nq` and echo it to stdout
- To store an output of a program in a variable, we write “`TEMPDIR=$(mkdir -p /tmp/${USER}/XXXXXXXX)`”

Let's write a little Shell script...

## Section 3

### Specifics of GWGD HPC cluster

## 2 filesystems

- ① HOME filesystem
- ② SCRATCH filesystem

### HOME

- Stores your *permanent* data.
- There is a quota. It could be extended on request.
- Has a backup mechanism.

### SCRATCH

- Stores your *temporal* data used for computations or projects.
- Fast and large filesystem.
- No Quota, but there are some rules to use it.

## HOME

- Quota is set per user basis.
- Find it out using Quota command

```
gwdu101:14 11:55:41 ~ > Quota
```

```
Global Filesystem KBytes: used softlimit hardlimit ...
UNI11                370216          0          0
UNI05                65316256 104857600 419430400
```

## SCRATCH

- No Quota per user. However, storage is limited

```
gwdu101:45 10:52:46 ~ > df -h /scratch
```

```
Filesystem      Size  Used Avail Use% Mounted on
beegfs_nodev   1.6P  750T  796T  49% /scratch1
```

- **local** filesystem is NOT shared, but fast (SSDs).
- Use it for temporal data on every node
- At the moment the ampXXX nodes have only a ramdisk as /local (very fast, but small, as it is part of your memory allocation).
- The size of it rather small

```
bash-4.2$ df -h /local
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sda6	78G	57M	74G	1%	/local



## Archive location

- Personal archive is located at `/usr/users/a/USERNAME`
- You can get the path from `$AHOME` variable

## Usage

- It is necessary to compress directories as tar or zip files
- if you want to archive directory data, call

```
tar -czvf $AHOME/data.tgz data
```

or faster (uses 4 cores and faster compression)

```
PIGZ="-1 -p 4 -R" tar -I pigz -cvf $AHOME/data.tgz data
```

## Important

The Scratch filesystem is **NOT** a permanent storage

## Recommended workflow

- Create directory for your project  
/scratch/users/\$USER/PROJECT
- Copy all necessary data there
- Run your compute jobs
- After completion of your jobs, save important results, that you need for further work to your home directory
- Delete all temporary files and broken runs
- Move the rest of the directory, that you want to keep for reference, into the archive and delete it from Scratch

```
tar -czvf $AHOME/PRJ.tar.xz /scratch/users/$USER/PROJECT
rm -rf /scratch/users/$USER/PROJECT
```

Try it!



- Create a project directory for this course:

```
mkdir /scratch/users/${USER}/scc-course
```

- Add some files in it

```
echo "a" > /scratch/users/${USER}/scc-course/file1
```

```
echo "b" > /scratch/users/${USER}/scc-course/file2
```

- Compress the folder and send to archive

```
tar -czvf $HOME/scc-course.tar.xz /scratch/users/${USER}/scc-course
```

There are 2 transfer servers that can be used to transfer data from your machine to HPC.

[transfer.gwdg.de](https://transfer.gwdg.de)

- reachable from the Internet
- only HOME is mounted

[transfer-scc.gwdg.de](https://transfer-scc.gwdg.de) (new name [transfer-mdc.hpc.gwdg.de](https://transfer-mdc.hpc.gwdg.de))

- reachable only from GÖNET
- HOME and /scratch are available

## SCP

*works on Linux, macOS, and latest Windows*

```
scp -rp {SRC-DIR} {USER}@transfer.gwdg.de:{DST-DIR}
```

to transfer back, simply swap the arguments

```
scp -rp {USER}@transfer.gwdg.de:{SRC-DIR} {DST-DIR}
```

## Filezilla

*works on all platforms. GUI. Open source software.*

## Rsync

*works on Linux, macOS*

```
rsync -avvH {SRC-DIR} {USER}@transfer.gwdg.de:{DST-DIR}
```

to transfer back, simply swap the arguments

```
rsync -avvH {USER}@transfer.gwdg.de:{SRC-DIR} {DST-DIR}
```

# Screen. Resuming interactive sessions



Screen – is the utility which allows you to resume the sessions.

## Usage

`screen` starts a screen session

`screen -S SName` starts a named screen session

`screen -r SName` resume the screen SName

`screen -ls` list all your available screens

within the screen you work as in usual shell

all screen commands start with `Ctrl + a`

`Ctrl + a d` detach from a screen session

`Ctrl + a c` create a new window

`Ctrl + a 0` switch to window 0, or use another number

## Section 4

### Preparing the environment with “modules”

- “module avail” find a list of installed modules
- “module list” list of currently loaded modules
- “module load software/version”
- “module purge” unload all modules
- “module unload software” unload a single module
- Most of the modules just append or prepend a path to PATH and MANPATH variables.
- Or set default variables to be found by compiler/configure scripts at compile time.



- Software provided as modules are compiled for specific CPU architecture: Cascadelake or Haswell.
- Names of these modules are the same, the correct version is loaded depending on the node you(your jobs) are.
- If you compile your software for specific architecture, check the modules you are using with `module whatis` command. It contains the "Target".

```
> gwdu103 ~ > module whatis gromacs
> ...
> gromacs/2020.4      : Target : haswell
> gwdu101 ~ > module whatis gromacs
> ...
> gromacs/2020.4      : Target : cascadelake
```

## Section 5

# Compiling Software

# Why Compiling?



- Compiling means to create an executable – or a library – from the source code
- GWWDG cannot install all software required by users (see modules for what is available)
- Scientific software is often only available as source code
- Compiling on the target system often yields better performance
- Prepackaged software typically requires administrator (root) privileges ...
  - ➔ (sudo or su won't work)
  - ➔ but you can use Singularity containers!

Singularity is the containerization system, just like Docker.  
However, we don't provide Docker in HPC for security reasons.

## Usage

To load singularity use the modules

```
module load singularity/3.7.0
```

You can run either native Singularity or Docker images.

```
singularity run library://sylabscsd/examples/lolcow
```

With Docker image

```
singularity run docker://godlovedc/lolcow
```

Some software packages provide Docker or Singularity images, if they do it will be easier to run them as containers.

Try it!

- Source code is usually packaged as “tarball”
  - Look for file extensions “tar.gz”, “tar.bz2”, “tgz”
  - Naming convention is often {NAME}-{VERSION}.tar.gz
- If the tarball is available on the web use “wget” to download
- Use “tar” to unpack the tarball
  - Use “tar xvzf” for “tar.gz”, “tgz”
  - Use “tar xvjf” for “tar.bz2”

## Using wget and tar to prepare the source code

```
> mkdir $HOME/build  
> cd $HOME/build  
> wget <tarball URL>  
> tar xvzf <name-version>.tar.gz  
> cd <name-version>
```

- Standard method: “./configure; make; [make check; make install]”
- Without root privileges: “--prefix” at configuration
- For better performance: Use Intel compilers and MKL
- For MPI (distributed parallel) applications: Use Intel MPI

## About “--prefix”



- “--prefix” is used to specify the base directory for your software
- use “./configure --prefix=DIR” to install directly in DIR.
- e.g. “./configure --prefix=\$HOME/software/<name-version>” to install into a software specific directory.



## Building and installing software into a specific directory

```
> cd $HOME; mkdir software
> cd $HOME/build/<name-version>
> ./configure --prefix=$HOME/software/<name-version>
> make -j 4; make check
> make install
> ln -s $HOME/software/<name-version>/bin/* $HOME/bin
> ln -s $HOME/software/<name-version>/lib/* $HOME/lib
> ln -s $HOME/software/<name-version>/include/* $HOME/include
```

- The GNU compilers (`gcc`, `gfortran`) are the standard compilers in Linux
- Other compilers are often faster, especially for Fortran code
- Recommended for overall performance: Intel compilers (`icc`, `ifort`)

## Building and installing software with Intel compilers

```
> module load intel
> CC=icc; CXX=icpc; FC=ifort; F77=ifort; F90=ifort
> export CC CXX FC F77 F90
> ./configure --prefix=$HOME/software/<name-version>
> make -j 4; make check
> make install
```

- A (shared) library is a collection of thematically related subroutines ready to use in a program
- The process of connecting a library to the (compiled) program is called linking
- Intel's Math Kernel Library provides performance optimized linear algebra and Fourier transform functions

## Example: linking programs to MKL

```
> module load intel
> CC=icc; CXX=icpc; FC=ifort; F77=ifort; F90=ifort
> export CC CXX FC F77 F90
> module load intel-parallel-studio
> export CPPFLAGS="-I${MKLROOT}/include -I${MKLROOT}/include/fftw"
> export LDFLAGS="-L${MKLROOT}/lib/intel64 -lmkl_intel_lp64\
> -lmkl_sequential -lmkl_core -lpthread -lm"
> ./configure --prefix=$HOME/software/<name-version>
> make -j 4; make check
> make install
```

Use Intel MKL Link Line Advisor!

<https://software.intel.com/en-us/articles/intel-mkl-link-line-advisor>

- MPI programs are meant to run distributed across several computers
- They require to be linked to an MPI library
- The recommended MPI library at GWGD is Intel MPI
- Another option is OpenMPI

## Building MPI programs with Intel MPI

```
> module load intel
> module load intel-mpi
> CC=mpiicc; CXX=mpiicpc; FC=mpiifort; F77=mpiifort; F90=mpiifort
> export CC CXX FC F77 F90
> module load intel-parallel-studio
> export CPPFLAGS="-I${MKLRROOT}/include -I${MKLRROOT}/include/fftw"
> export LDFLAGS="-L${MKLRROOT}/lib/intel64 -lmkl_intel_lp64\
> -lmkl_sequential -lmkl_core -lpthread -lm"
> ./configure --prefix=$HOME/software/<name-version>
> make -j 4; make check
> make install
```

## Preparation

```
> module load openmpi
> echo $OPENMPI_ROOT
/opt/sw/rev/20.12/cascadelake/gcc-9.3.0/openmpi-4.0.5-ub54mm
> R
```

## R command line

```
> install.packages("Rmpi", dependencies=TRUE,
  configure.args=c("--with-mpi=/opt/sw/rev/20.12/cascadelake/gcc-9.3.0/
  openmpi-4.0.5-ub54mm"
  ))
> install.packages(c("foreach", "doMPI"))
```



Spack is a package management tool designed to support multiple versions and configurations of software on a wide variety of platforms and environments.

Most of the software is installed using Spack and can be loaded either by using **module** or **spack-user** (Spack installation for users).

Spack-user installation allows to compile software yourself:

- Load spack-user module with  
`module load spack-user`
- Follow the instructions to activate shell support  
`source $SPACK_USER_ROOT/share/spack/setup-env.sh`
- Install supported software  
`spack install SPEC`
- Load the installed software  
`spack load SPEC`

# Spack install example



```
> module load spack-user
> source $SPACK_USER_ROOT/share/spack/setup-env.sh
> spack install ffmpeg@4.4
==> Warning: Missing a source id for ffmpeg@4.4
[+] /opt/sw/rev/20.12/haswell/gcc-9.3.0/alsa-lib-1.2.3.2-v63tdf
[+] /opt/sw/rev/20.12/haswell/gcc-9.3.0/libiconv-1.16-s5eu5n
[+] /opt/sw/rev/20.12/haswell/gcc-9.3.0/yasm-1.3.0-2fd743
[+] /opt/sw/rev/20.12/haswell/gcc-9.3.0/zlib-1.2.11-f2qirj
[+] /opt/sw/rev/20.12/haswell/gcc-9.3.0/diffutils-3.7-bv5dar
[+] /opt/sw/rev/20.12/haswell/gcc-9.3.0/bzip2-1.0.8-fdvhu3
==> Installing ffmpeg-4.4-3jhy2clncz2upmcthclpu436ld5qsa22
==> Warning: There is no checksum on file to fetch ffmpeg@4.4 safely.
==> Fetch anyway? [y/N] y
==> Fetching http://ffmpeg.org/releases/ffmpeg-4.4.tar.bz2
==> ffmpeg: Executing phase: 'autoreconf'
==> ffmpeg: Executing phase: 'configure'
==> ffmpeg: Executing phase: 'build'
==> ffmpeg: Executing phase: 'install'
[+] /usr/users/akhuziy/.spack/install/haswell/gcc-9.3.0/ffmpeg-4.4-3jhy2c
> spack load ffmpeg@4.4
> ffmpeg --version
ffmpeg version 4.4 Copyright (c) 2000-2021 the FFmpeg developers
```

# Table of Contents, Part II



- 1 Connecting to the frontends
- 2 The most important Linux commands
- 3 Specifics of GWDG HPC cluster
- 4 Preparing the environment with “modules”
- 5 Compiling Software
- 6 Using Slurm - Basics
- 7 Using Slurm - Advanced
- 8 Getting Help

## Section 6

### Using Slurm - Basics

#### Getting started with Slurm

- Cluster divided into frontends and compute nodes
- Compute nodes to all calculations
- You cannot connect directly to the nodes
- You cannot run heavy calculations on the frontends
- So how do you use the compute nodes?

Use our scheduler: Slurm


# How to use the cluster



- Cluster divided into frontends and compute nodes
- Compute nodes to all calculations
- You cannot connect directly to the nodes
- You cannot run heavy calculations on the frontends
- So how do you use the compute nodes?

Use our scheduler: **Slurm!**

# How to use the Cluster





A job is a set of instructions for Slurm, including

- one or multiple programs to execute
- estimated runtime
- required resources (CPUs, GPUs, Memory)
- and more...

Use `srun` to submit a job to Slurm

```
srun <program>
```

Example:

```
gwdu101:27 12:53:50 ~ > hostname
```

```
gwdu101
```

```
gwdu101:27 12:53:53 ~ > srun hostname
```

```
amp078
```

```
gwdu101:27 12:53:56 ~ > srun hostname -f
```

```
amp078.global.gwdg.cluster
```

- `srun` submits information on your job to Slurm
  - What is to be done? (path to your program and required parameters)
  - What are its requirements? (e.g. which nodes, number of tasks, maximum runtime)
- Slurm matches the jobs requirements against the capabilities of our nodes
- When suitable free resources are found the job is started
- Slurm prioritizes the jobs based on a number of factors.

- Different compute nodes have different features
- Slurm differentiates using **Partitions**

General purpose partitions:

**medium** General purpose partition, well suited for most jobs. Up to 1024 cores per job.

**fat** Up to 512 GB in one host.

**fat+** For extreme memory requirements. Up to 2048GB per host.

Special purpose partitions:

**gpu** For jobs using GPU acceleration.

**int** For interactive jobs, i.e. jobs which require a shell or a GUI.

**Cluster** A collection of networked computers intended to provide compute capabilities.

**Node** One of these computers, also called host or server.

**Frontend** Special node provided to interact with the cluster. `login-mdc|fas.hpc.gwdg.de` in our case.

**Job** Program consisting of one or several parallel tasks.

**Partition** A pool of nodes suitable for the job

**Batch System** Management system distributing job tasks across job slots. Slurm.

```
srun <parameters> <program>
```

## common parameters

- p <partition> partition.
- t <hh:mm:ss> Maximum runtime. If this is exceeded the job is killed.
- A <all> Specify account 'all'. Only necessary if you are part of a working group with access to special partitions.

## `srun`: Interactive jobs

- `--x11` Adds X11 (GUI) forwarding. This requires that you connect to the frontend with `ssh -Y` and your local machine supports X-Windows.
- `-p int` Use the interactive partition. In `int` the nodes have no slot limit. They will take jobs until their load crosses a specified threshold, so jobs start immediately.
- `--pty` interactive mode



## Running Matlab

```
> ssh -Y login-mdc.gwdg.de  
> module load matlab  
> srun --x11 -p medium matlab
```

- The job will be dispatched and as soon as an available node is found and the Matlab interface will start.
- If you have your own license for Matlab then you need to place your `license.lic` file in `$HOME/.matlab/R2015a_licenses` directory (depending on the version you are using).

## Running R interactively

```
> ssh login-fas.hpc.gwdg.de  
> srun --pty -p medium R
```

Try it!

- Serial job** Job consisting of one task using one job slot.
- SMP job** Job with shared memory parallelization (often realized with OpenMP), meaning that all processes need access to the memory of the same node. Therefore uses several job slots **on the same node**.
- MPI job** Job with distributed memory parallelization, realized with MPI. Can use several job slots on several nodes and needs to be started with a helper program, e.g., `mpirun` or `srun`.

- MPI jobs are a lot of independent tasks that (usually) use one core each
  - ➔ started with `srun` or `mpirun`
  - ➔ Slurm calls these tasks
- Single node jobs are usually just one task with many cores
- Both can be combined into hybrid jobs: multiple tasks using multiple cores each

## srun options for parallel (SMP or MPI) jobs.

- N <min>-<max>, --nodes=<min>-<max> Minimum and maximum node count. You can also specify the exact number.
- n,--ntasks=<n> Number of tasks (not equally distributed!)
- tasks-per-node=<n> Tasks per node. If used with -n it denotes the maximum number of tasks per node.
- c,--cpus-per-task=<n> CPUs per tasks.

## Rule of thumb

- -c for single node jobs
- -n for MPI jobs

## Rule of thumb 2

If you are unsure if your program uses MPI, then it does not.

# A note on -n vs. -c



## Rule of thumb

- -c for single node jobs
- -n for MPI jobs

## Rule of thumb 2

If you are unsure if your program uses MPI, then it does not.



Reservation: `scc-course`

## Usage

Either: `--reservation=scc-course` for each job

Or: `export SLURM_RESERVATION=scc-course`

The latter has to be unset, if you want to submit to a partition besides medium.

Try it!



## Exercises

Try these job configurations

- 1 10 tasks
- 2 10 tasks distributed over 3 nodes
- 3 3 nodes with 3 tasks each
- 4 1 task with 5 cores
- 5 2 tasks per node on 2 nodes with 4 cores per task

use `slurm_resources` script to get see the resources of your job

## srun options

`--mem <size[K|M|G|T] >` Memory per node.

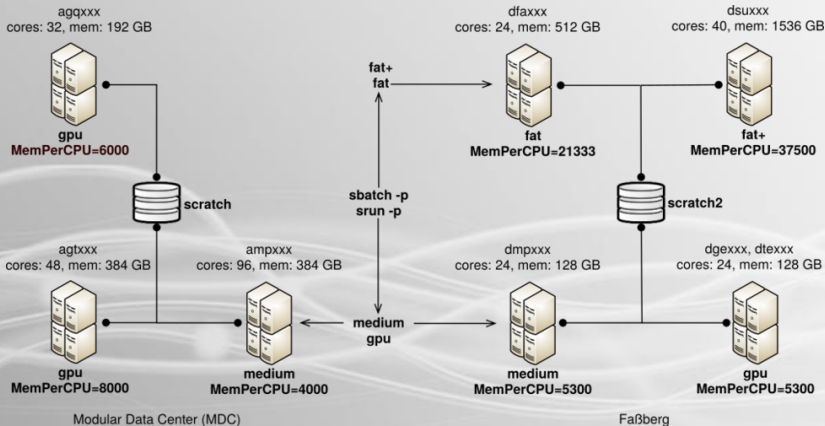
`--mem-per-cpu <size[K|M|G|T] >` Memory per core.

- without options:

- each partition has a `DefMemPerCPU` option

- can be retrieved via `scontrol show partition <name>`

# The GWDG Scientific Compute Cluster



Try it!



### Exercise:

Play with the combination of number of cores or tasks, nodes and their effect on your available memory:

- 1 1 core and `--mem 4G`
- 2 3 tasks and 2 nodes, see effect of `--mem` and `--mem-per-cpu`
- 3 20 tasks, see distribution of memory over hosts.

## Problem

- if you have big jobs, your queue time will be long
- `srun` needs you to stay logged in
- jobs can run for days

## Solution

Batch Jobs!

## Problem

- if you have big jobs, your queue time will be long
- `srun` needs you to stay logged in
- jobs can run for days

## Solution

Batch Jobs!

# sbatch: Using Job Scripts



A job script is a shell script with a special comment section.  
The #SBATCH lines have to come first!

## sbatch: Basic job script example

```
#!/bin/bash
#SBATCH -p medium
#SBATCH -t 10:00
#SBATCH -o job-%J.out
```

slurm\_resources

Submit with:

```
sbatch <script name>
```



- A job script is essentially a normal script
- usually bash/shell, but can be any scripting language (R, python, perl)
- #SBATCH lines need to be at the top!
- you can copy files, load modules, and do any scripting you want
- for MPI, use `srun` or `mpirun` to start your program

## Solution

```
sbatch <slurm options> jobscript
```

- `--mail-type=<TYPE>` get mail notifications (type: BEGIN, END, etc.)
- `--mail-user=<address>` Default: `${USER}@gwdg.de`
- `-o/-e <file>` Store job output in file (slurm-`<jobid>.out` by default). `%J` in the filename stands for the jobid.

**sinfo** Info about the system and partitions.

`-p <partition>, -t <state>`

**squeue** Show the job queue.

`-p <partition>, --me`

**scancel** Cancel Job

`scancel <JobID>`

`scancel -p <partition>|-u $USER`

## Exercise

Write your own Job script.

- Use `echo`, `hostname`, and `sleep X` (sleep for X seconds) to generate output or have it running for a longer time.
- Have the job send you an email. Advanced: Take a look at the different mail-type options. What do they do?
- Write the output to a different file. Redirect output and error into different files. Advanced: Take a look at the filename pattern options. Include node and job name in the output file.

Time: 20 Minutes

## Distributing tasks in the medium partition

```
#SBATCH -p medium  
#SBATCH -n 240  
#SBATCH -o job-%J.out
```

```
module purge  
module load intel/compiler intel/mkl intel/mpi namd
```

```
srun namd2 +setcpuaffinity apoa1.namd
```

This will spread tasks among many nodes.

## Distributing tasks in the medium partition

```
#SBATCH -p medium
#SBATCH -N 10
#SBATCH --ntasks-per-node 24
#SBATCH -o job-%J.out

module purge
module load intel/compiler intel/mkl intel/mpi namd

srun namd2 +setcpuaffinity apoa1.namd
```

Memory is faster than network!

Try to spread your tasks to as little nodes as possible.


**/local** Local hard disk of the node. SSD based and therefore a very fast option for storing temporary data. Automatic file deletion. A temporary directory is created on all nodes at `$TMP_LOCAL`.

**/scratch** Shared scratch space, available on most nodes, but there are two instances (use `-C scratch` or `-C scratch2`). Very fast, no automatic file deletion, but also no backup! Files may have to be deleted manually when we run out of space.

**/scratch/ssd** special extra fast scratch file system only on `scratch1`. Ideal for temporary data in jobs spanning multiple nodes. Automatic file deletion. A per-job directory is created at `$TMP_SCRATCH`.

**\$HOME** Available everywhere, permanent, with backup. Personal disk space can be increased. Comparably slow.

# The GWDG Scientific Compute Cluster





# Recipe: Using /scratch



```
#!/bin/bash
#SBATCH -p medium
#SBATCH -n 24
#SBATCH -N 1
#SBATCH -C scratch
#SBATCH -t 1-00:00:00

export g09root="/usr/product/gaussian/g09/d01"
source $g09root/g09/bsd/g09.profile

if [ ${TMP_SCRATCH} -a -d ${TMP_SCRATCH} ]; then
    export GAUSS_SCRDIR=${TMP_SCRATCH}
else
    export GAUSS_SCRDIR=$TMP_LOCAL
fi

g09 myjob.com myjob.log
```

# Try it!



## Exercise

Write a job script, where you

- create a scratch directory
- copy data from your home file system to the scratch directory
- run a job with the data
- copy the results back
- delete the scratch directory

If you do not have a program/data to try this on, there is a small python program in `/scratch1/projects/scc-course/` and a bit of input data.

# The fat+ partition



The fat+ partition contains:

- 5 nodes with 1.5Tb Memory
- 1 node with 2Tb Memory

Usage recommendations:

- Work your way up. Start in fat and only use fat+ if your jobs runs out of memory.
- Use `sacct` or `profit-hpc`, see if your job really is memory bound
- When unsure, ask us!
- `--memor --mem-per-cpu` is mandatory
- You might get angry mails from me, if you waste resources here

## Running hybrid jobs

```
#SBATCH -p medium
#SBATCH -N 5
#SBATCH --ntasks-per-node=4
#SBATCH --cpus-per-task=6
#SBATCH -o job-%J.out

module purge
module load openmpi/gcc

export OMP_NUM_THREADS=$SLURM_CPUS_PER_TASK

srun hybrid_job
```

## The `--qos` parameter

- Default maximum runtime: 2 days
- `--qos= <qos>` can select a QoS
- Two extra QoS available:
  - `short` for shorter jobs (max. 2h), has higher priority, limited job slots
  - `long` longer jobs (max. 5d), limited job slots.

## But my job is even longer

- try parallelizing more
- break it down into smaller steps
- check, if your software supports checkpoints
- check again!
- contact us

`scontrol show [partition|node|job] <x>` where x should be a node name, JobID or partition name.

`sprio` Priority information about pending jobs

`sacct` Get information about a job after it finished

`-j <jobid>`

`--format=JobID,User,JobName,MaxRSS,Elapsed,Timelimit`

## GPU parameters

`-G | --gpus=[type:]<n>` requests  $n$  GPUs of type

`--gpus-per-task=[type:]<n>` requests  $n$  GPUs of type per task

`--gpus-per-node=[type:]<n>` requests  $n$  GPUs of type per node

- CPUs are evenly distributed for every GPU

- Available types are:

- ➔ `rtx5000`

- ➔ `v100`

- ➔ `k40`

- ➔ `gtx1080`

- ➔ `gtx980`

- See: `sinfo -p gpu --format=%N,%G`

- take a look at your output files, while the job is running:
  - ➔ `tail -f /path/to/output`
- take a look at the jobs, while it is running
  - ➔ you can `ssh` into every node that currently calculates your job
  - ➔ use `htop` to see the processor and ram usage



## Read the extra job information

```
=====  
JobID = 4383174  
User = mboden, Account = admin  
Partition = gpu, Nodelist = dge[001,006]  
=====
```

```
[job output]
```

```
===== Job Information =====  
Submitted: 2020-04-24T17:35:41  
Started: 2020-04-24T17:35:41  
Ended: 2020-04-24T17:45:45  
Elapsed: 10 min, Limit: 60 min, Difference: 50 min  
CPUs: 2, Nodes: 2
```

```
===== Profit-HPC =====  
To generate the Profit-HPC text report, run the following command  
profit-hpc 4383174  
=====
```

Take a look at all the information. Is it as expected?

## Read your errors!

```
slurmstepd: error: Detected 1064 oom-kill event(s) in step XXXXXX.0 cgroup.  
Some of your processes may have been killed by the cgroup out-of-memory handler.  
srun: error: gwda024: task 3: Out Of Memory
```

Might have something to do with memory!

Have a look at your jobs memory with:

```
sacct -j JOBID -o jobid,MaxRSS,MaxRSSNode
```

And for more advanced job statistics, use profit-hpc

## Section 7

### Using Slurm - Advanced

Even more possibilities!

- For 'embarrassingly' parallel jobs
- starts an array of very similar jobs from a blueprint
- Please DO NOT use for loops for this

# Simple Example



```
#!/bin/bash
#SBATCH -p medium
#SBATCH -t 10:00
#SBATCH -n 1
#SBATCH -c 4

module load python/3.9.0
for i in {1..100}; do
    python myprogram.py $i
done
```

Spawns one job with 100 iterations sequentially

```
#!/bin/bash
#SBATCH -p medium
#SBATCH -t 10:00
#SBATCH -n 1
#SBATCH -c 4
#SBATCH -a 1-100
```

```
module load python/3.9.0
python myprogram.py $SLURM_ARRAY_TASK_ID
```

Spawns 100 jobs in parallel with 1 iteration.

---

<b>Notation</b>	<b>Effect</b>
-a 0-5	Range, from 0 to 5
-a 1,5,12	Specific Numbers
-a 0-5:2	Step size 2: in 0, 2, 4
-a 0-11%4	Only four jobs at once.
-a 0-11%1	Run all jobs sequentially
-a 0-5:2,20%2	Everything combined. 0,2,4,20, but only two at a time.

---

- control jobs via environment variables:
  - ➔ `SSLURM_ARRAY_TASK_ID`
  - ➔ `SSLURM_ARRAY_TASK_COUNT`
  - ➔ `SSLURM_ARRAY_TASK_MAX`
  - ➔ `SSLURM_ARRAY_TASK_MIN`
  - ➔ `SSLURM_ARRAY_TASK_STEP`



## More Complex Example



```
#!/bin/bash
#SBATCH -p medium
#SBATCH -t 10:00
# insert X as the number of .inp files you have -1
# ls *.inp | wc -l
#SBATCH -a 0-X
```

```
shopt -s nullglob
arr=(./*.inp)
```

```
./big_computation ${arr[$SLURM_ARRAY_TASK_ID]}
```

Or leave out the -a line and add to sbatch

```
sbatch -a 0-$(($(ls ./*.inp | wc -l)-1)) jobarray.sh
```

Wait for a specific job to finish, before the next starts:

```
-d | --dependency=dependency_definition
```

where `dependency_definition` can be:

`after:job_id[+time]` After the specified jobs start or are cancelled

`afterok:job_id` After the specified jobs have successfully executed

`afternotok:job_id` After the specified jobs have terminated in some failed state

`afterany:job_id` After the specified jobs have terminated.

- use `sbatch --parsable` when scripting dependencies.
- If dependency is not met, job needs to be canceled manually
- More details in the `sbatch` man-page or online help

# Miscellaneous Slurm Options



`--wrap=` wrap the specified command string in a simple "sh" shell script. Only for `sbatch`.

`--test-only` Check script and give estimate when it would run.

`--open-mode=append|truncate` append or overwrite job files

`--export=NONE` don't export user environment, helpful for reproducibility.

`--signal=B:12@600` Send signal 12 to job when 600 seconds before time limit. You can catch the signal in the script:

```
[...]  
trap 'cp -af ${TMP_LOCAL}/* /scratch/your_dir/; exit 12' 12  
your_job &  
wait
```

- use job arrays where possible (don't sbatch in a for loop)
- set a reasonable time limit (not just 2 days)
- use the short QoS where applicable
- ask us!

## Section 8

### Getting Help

- man pages
- Slurm online help
  - ➔ For example: `sbatch --help`
- GWDG scientific compute cluster documentation
  - ➔ [https://docs.gwdg.de/doku.php?id=en:services:application\\_services:high\\_performance\\_computing:start](https://docs.gwdg.de/doku.php?id=en:services:application_services:high_performance_computing:start)
- GWDG scientific compute cluster user wiki
  - ➔ <https://info.gwdg.de/wiki/doku.php?id=wiki:hpc:start>
- HPC announce mailing list
  - ➔ <https://listserv.gwdg.de/mailman/listinfo/hpc-announce>

- Everyone with a cluster account can add to the Wiki!
- Please inform us of all changes and new articles at [parallel@gwdg.de](mailto:parallel@gwdg.de).
- Please add the category "*Scientific Computing*" to all contributions regarding the cluster.



- Rocket.Chat Channel at:  
<https://chat.gwdg.de/channel/hpc-users>
- Mainly for quick questions
- Usually faster response but we might tell you to open a ticket

- Write an email to *hpc-support@gwdg.de*
- State your user id (`$USER`)
- If you have a problem with jobs, **always** include:
  - Job IDs
  - standard output ( `-o <file>`)
  - standard error ( `-e <file>`)
- If you have a lot of failed jobs send at least two outputs. You may also list the jobid's of all failed jobs.
- If you don't mind us looking at your files, please state this in your request
  - You may limit your permission to specific directories or files

- Bi-weekly meeting between users and admins.
- Format: Short focus session and open discussions and Q&A afterwards
- More Information at [https://docs.gwdg.de/doku.php?id=en:services:application\\_services:high\\_performance\\_computing:hpc\\_coffee](https://docs.gwdg.de/doku.php?id=en:services:application_services:high_performance_computing:hpc_coffee)

The End



Thank you all for your attention!

Questions?

- Convention: Executables are stored in “bin”, shared libraries in “lib” directories
- Directories in “\$PATH” are searched for binaries, directories in “\$LD\_LIBRARY\_PATH” for libraries
- Two strategies:
  - ① Put everything directly under \$HOME/bin, \$HOME/lib
    - Easy to setup search paths
    - Difficult to remove software packages
  - ② Put each software into its own subdirectory
    - Easy to remove software (with “rm -rf <subdirectory>”)
    - Difficult to setup search paths

- Or combine both strategies:
  - ➔ Put each software in its own subdirectory
  - ➔ Use “`ln -s`” to link everything to `$HOME/bin` and `$HOME/lib`, respectively
  - ➔ Use “`export LD_LIBRARY_PATH=$HOME/lib:$LD_LIBRARY_PATH; export PATH=$HOME/bin:$PATH`” in your shell and scripts
  - ➔ Use “`find $HOME/bin $HOME/lib -xtype l -delete`” after removing software