

DAY 1: SUPERCOMPUTING

Getting started on a supercomputer

2021-02-01 | Stefan Kesselheim | Helmholtz AI @ JSC

OUTLINE

Supercomputers

How to get in

Queueing Jobs

Tutorial Tasks

WHAT IS A SUPERCOMPUTER


Compute cluster	Many computers bound together locally
Supercomputer	A damn lot of computers bound together locally
Grid computing	Many clusters working as some kind of supercomputer
Cloud computing	Virtual machines in compute center(s)
High throughput computing	Many independent computations



ANATOMY OF A SUPERCOMPUTER

- Login Nodes: Normal machines, for compilation, data transfer, scripting, etc. No GPUs.
- Compute Nodes: Guess what :-)
- Special Purpose Nodes: Visualization, ...
- Network file system
- Scratch file system accessible from compute nodes
- Key stats:
 - Number of Nodes
 - CPUs, Number of Cores, Single-core Performance
 - RAM
 - Network: Bandwidth, Latency
 - Accelerators

Structure



SUPERCOMPUTER@JSC

JURECA DC

432 Nodes, Dual Socket AMD EPYC with 64 cores@ 2.25 GHz, 512/1025 GiB RAM, InfiniBand HDR 100, 48 nodes with 4x Nvidia A100 with 40gb.

JUWELS

2271+240 Nodes, Dual Socket Skylake with 24 cores@ 2.7 GHz, 96, 196 GiB memory, InfiniBand EDR. 56 Nodes with 4x NVidia V100 16gb.

JUWELS Booster

936 Nodes, 2x AMD EPYC Rome 7402 CPU, 2x 24 cores, 2.7 GHz, 512, 1024 GiB memory, Mellanox HDR, 4x NVIDIA A100 with 40gb

DEEP-EST

Modular Prototype. Cluster Module, 50x2x Xeon 12core @3.2 GHz. Extreme Scale Booster, 75x Xeon 8Core@2.5 GHz w 1x Nvidia V100 (compact!). Data Analytics Module with 1xNvidia V100 and 1x Intel FPGA Stratix10.

JUST



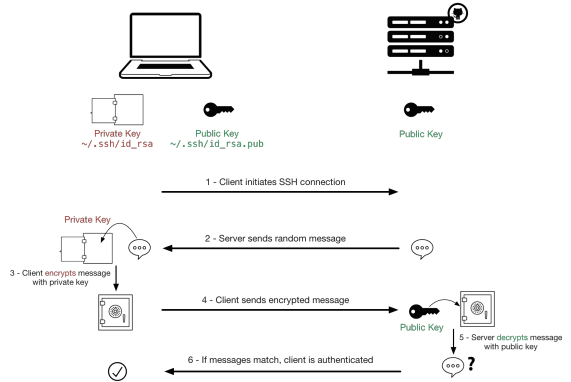
- 52 PB total storage
- \$HOME: Minimal, only dotfiles
- \$SCRATCH: 9 PB, Working storage, deleted after 90 days.
- \$DATA: 14 PB, Large scale file storage, not available from compute nodes.
- \$FASTDATA: 9 PB, Large scale file storage that must be continuously accessed.
- \$PROJECT: 2.3 PB, Project data .
- Peak bandwidth up to 400 GB/sec.
- Data is shared among all supercomputers.

AUTHENTICATION AND AUTHORIZATION

- Create public/private key pair
- Register public key in JuDoor
- Make sure to use that key.

```
ssh -i path/to/keyfile  
surname1@juwels.fz-juelich.de
```

```
kessel@DESKTOP-LM2SDL3: /mnt/c/Windows/system32$ ssh kesselheim1@juwels.fz-juelich.de  
Last login: Mon Jan 25 20:37:13 2021 from ip-88-153-87-68.hsi04.unimediagroup.de  
  
*****  
Welcome to                               *  
JUWELS                                   Juelich Wizard *  
                                           for             *  
                                           European Leadership *  
                                           Science           *  
*****  
2020-11-19T14:00+0200
```



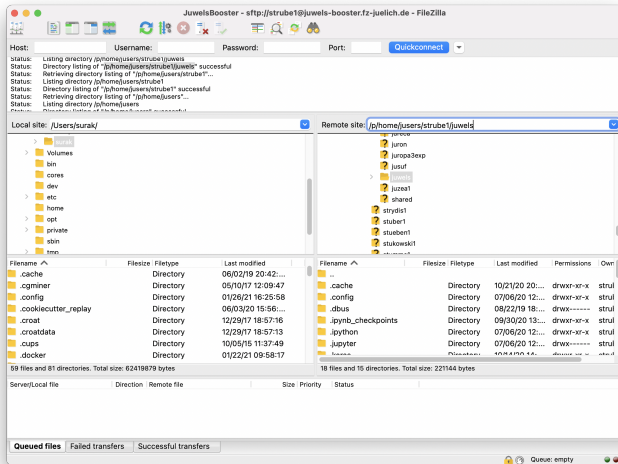
Attention: You will need to restrict IP ranges/DNS areas. Check tutorial.

FILE ACCESS

Copying files to a remote machine

`scp -i path/to/key <source> <target>`

Or use a graphical tool, e.g. Filezilla:



FILE ACCESS (CONT.)

Your home directory is **tiny**.

Project directory: /p/project/training2004

```
cd /p/project/training2004
mkdir $USER
cd $USER
git clone https://gitlab.version.fz-juelich.de/MLDL_FZJ/juhaicu/jsc_internal/superhaicu/
shared_space/teaching/intro_scalable_dl_2021/course-material
```

SUPERCOMPUTER USAGE MODEL

- Using the the supercomputer means submitting a job to a batch system.
- Job scheduling according to priorities. The jobs with the highest priorities will be scheduled next.
- No node-sharing. The smallest allocation for jobs is one compute node (4 GPUs).
- Compute time allocation is based on compute projects. For every compute job, a compute project pays.
- Data projects allocate large amounts of storage, but no compute.
- To enable fair share, only relatively short job runs (24h) are allowed. Please implement checkpointing (or make your code fast enough).
- Solution for long-running tasks: Job arrays.

Query Project quotas:

```
jutil project show -project training2004  
q_cpuquota
```

SLURM I: EXAMPLE SUBMISSION SCRIPT

```
#!/bin/bash
#SBATCH --nodes=1           # How many nodes?
#SBATCH -A training2004     # Who pays for it?
#SBATCH --partition booster # Where does the code run?
#SBATCH --gres gpu          # Not required on booster
#SBATCH --time=00:15:00     # How long?
#SBATCH -o output.txt
#SBATCH -e error.txt

source /p/project/training2004/course2021_working_environment/activate.sh
cd /p/project/training2004/${USER}
srun --ntasks-per-node=4 python my_script.py
```

sbatch my_script.sh

SLURM II: RESOURCES, COMPUTE BUDGET

Partition

Every node belongs to at least one partition of the cluster.

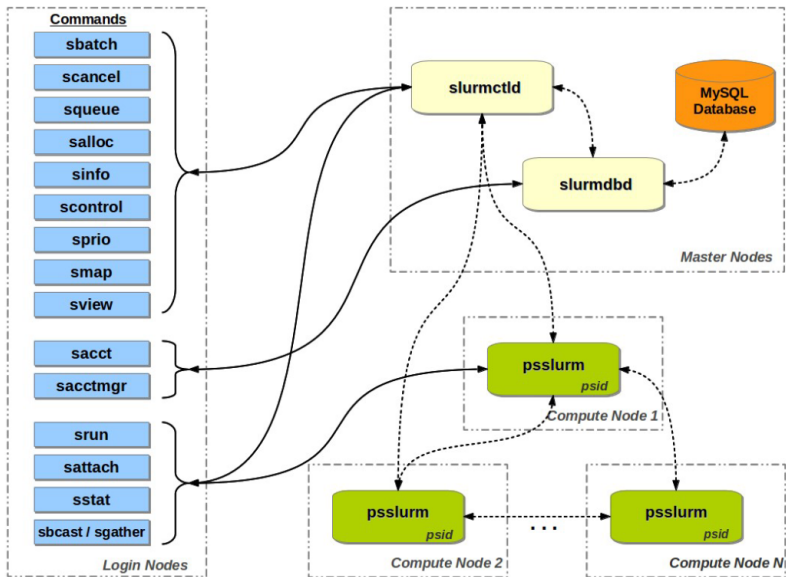
Queue

Jobs wait in a queue. Queues are executed in one or several partitions.

Generic Resources (GRES)

Jobs can request generic resources. One option is GPUs. RAM requirements or software licenses are also possible

SLURM III



MODULE SYSTEM

All kind of software is already installed in modules.

Compiler/GCCcore/9.3.0		
Autotools/20200321	Night-Compute/2020.3.0	h5py/2.10.0-serial-
Bazel/2.6.0	Night-Systems/2020.3.1	Python-3.8.5
CFITSIO/3.490	Night-Systems/2020.4.1	hwlac/2.2.0
CMake/3.18.0	Night-Systems/2020.5.1	jemalloc/5.2.1
CVS/1.11.23	Octave/6.1.0-nomei	libvex/1.9.0
Cling/0.7	OpenCV/4.5.0-Python-3.8.5	likwid/5.0.2
CubaGUJ/4.5	OpenEXR/2.5.2	mapama/2.5.4
Doxygen/1.8.18	OpenGL/2020	meld/2.21.0
Eigen/3.3.7	OpenJPEG/2.3.1	Python-3.8.5
Emacs/27.1	PAPI/6.0.0	memkind/1.10.1
FrBidi/1.0.9	PCRE2/10.34	nano/3.5
GDB/10.1	Par/5.32.0	netCDF-C++4/4.3.1-serial
GEOS/3.8.1-Python-3.8.5	Pillow-SIMD/7.0.0.post3-	netCDF-Fortran/4.5.3-
GMP/6.2.0	Python-3.8.5	serial
GPicView/0.2.5	PostgreSQL/12.3	netCDF/4.7.4-serial
GL/2.6	PyOpenCL/2020.2.2-Python-3.8.5	netcdf4-python/1.5.4-
Grace/5.1.25	PyQt5/5.15.1-Python-3.8.5	serial-Python-3.8.5
Graphviz/2.44.1	PyTorch/1.7.0-Python-3.8.5	numba/0.51.1-
HRF3/1.10.6-serial	Python-NeuroImaging/2020-	Python-3.8.5
ICA-AROMA/0.4.4-beta-	Python-3.8.5	parallel/20201122
Python-3.8.5	Python/3.8.5	pyproj/2.6.1.post1-
ITK/5.1.2-nomei-Python-3.8.5	R/4.0.2-nomei	Python-3.8.5
ImaceMapics/7.0.10-25	Ruby/2.7.1	re2c/1.3
Julia/CUDA/2.0.2	Rust/1.47.0	scikit/2020-Python-3.8.5
Julia/1.5.2	SciPy-Stack/2020-Python-3.8.5	tbp/2020.3
Jupyter/2020.2.5-Python-3.8.5	Shapely/1.7.1-Python-3.8.5	tsch/6.22.02
JupyterCollection/2020.2.5	Simulants-Tools/2020-	texlive/20200406
LLVM/10.0.1	Python-3.8.5	tmux/3.1b
MPFR/4.1.0	Subversion/1.14.0	torchvision/0.8.1-
Mercurial/5.5.2-Python-3.8.5	Tcl/8.6.10	Python-3.8.5
Meson/0.55.0-Python-3.8.5	TensorFlow/2.3.1-Python-3.8.5	trimesh/3.8.11-
NCCL/2.8.3-1-CUDA-11.0	TotalView/2020.1.13	Python-3.8.5
NSPR/4.23	VirtualGL/2.6.4	unzip/6.0
NSS/3.51	Voro++/0.4.6	xrca/4.0.4-Python-3.8.5
Ninja/1.10.0	X11/20200222	zsh/5.8
Night-Compute/2020.1.2	XServer/1.20.9	
Night-Compute/2020.2.0	ZeroMQ/4.3.3	

https://apps.fz-juelich.de/jsc/llview/juwels_modules_booster/

MODULE SYSTEM

Commands

- `module avail`
- `module purge # unload everything`
- `module load`
- `module spider`

`module spider nano`

```
kesselheim1@jwlogin07 ~]$ module spider nano
```

```
nano: nano/5.5
```

Description:

GNU nano is a small and friendly text editor. Besides basic text editing, nano offers features like undo/redo, syntax coloring, interactive search-and-replace, auto-indentation, line numbers, word completion, file locking, backup files, and internationalization support.

`module load nano`

For this Training:

```
source /p/project/training2004/course2021_working_environment/activate.sh
```

```
https://gitlab.version.fz-juelich.de/MLDL\_FZJ/juhaicu/jsc\_internal/superhaicu/  
shared\_space/teaching/intro\_scalable\_dl\_2021/course2021\_working\_environment
```

`module avail`

		Core packages	
Advisor/2020_update3		Python/3.8.5	
Autotools/20200321		R/4.0.2-nompi	
Autotools/20200321	(D)	Ruby/2.7.1	
Bazel/3.6.0		Rust/1.47.0	
Blender/2.90.1-binary		SciPy-Stack/2020-Python-3.8.5	
CFITSIO/3.490		Shapely/1.7.1-Python-3.8.5	
CMake/3.18.0		Singularity-Tools/2020-Python-3.8.5	
CUDA/11.0	(g)	StdEnv/2020	(L)
CVS/1.11.23		Subversion/1.14.0	
Cling/0.7		Tcl/8.6.10	
CubeGUI/4.5		TensorFlow/2.3.1-Python-3.8.5	
Doxygen/1.8.18		TotalView/2020.1.13	
EasyBuild/4.2.1		UCX/1.8.1	
EasyBuild/4.2.2		UCX/1.9.0	(D)
EasyBuild/4.3.0	(D)	VTune/2019_update8	
Eigen/3.3.7		Vampir/9.9.0	
Emacs/27.1		VirtualGL/2.6.4	
FriBidi/1.0.9		Voro++/0.4.6	
GDB/10.1		X11/20200222	

JUPYTER-JSC

<https://jupyter-jsc.fz-juelich.de/>

JupyterLab Options

System	JUWELS
Account	kesselheim1
Project	training2004
Partition	booster
Email notification	<input type="checkbox"/>
Nodes [1, 304]	1/3
Runtime (min) [1, 1440]	30/3
GPUs [1, 4]	4/3

Start

© Forschungszentrum Jülich Imprint Privacy Policy Support Terms of Service

HELMHOLTZ
RESEARCH FOR GRAND CHALLENGES

■ Add JupyterLab

- Machine (JUWELS)
- Compute Project (training2004)
- Partition (Login or Booster)
- Number of Nodes
- Runtime
- Start Jupyterlab
- Create Symlink: In a terminal, type
`ln -s /p/project/training2004/~ /training2004`
- Have fun

TUTORIAL TASKS

- 1 Connect to Juwels Booster by SSH
- 2 Run your first interactive Slurm Job
- 3 Queue your first batch job
- 4 Load the course compute environment
- 5 Bonus: Login into Jupyter-JSC