



New environment for CLAIX 2018

SLURM and new defaults for environment modules

SLURM as the batch system for CLAIX 2018

General differences between LSF and SLURM

- Submission
 - bsub < batchscript ⇔ sbatch batchscript
 - slurm NEEDS a batchscript, does not allow commandline submission

SLURM as the batch system for CLAIX 2018

General differences between LSF and SLURM

- Submission
 - `bsub < batchscript` ⇔ `sbatch batchscript`
 - slurm NEEDS a batchscript, does not allow commandline submission
- Magic cookie
 - `#BSUB` ⇔ `#SBATCH`
 - `#BSUB` cookies partly misinterpreted by SLURM → NOT allowed, prohibited by job modifier script

SLURM as the batch system for CLAIX 2018

General differences between LSF and SLURM

- Submission
 - bsub < batchscript ⇔ sbatch batchscript
 - slurm NEEDS a batchscript, does not allow commandline submission
- Magic cookie
 - #BSUB ⇔ #SBATCH
 - #BSUB cookies partly misinterpreted by SLURM → NOT allowed, prohibited by job modifier script
- Pending/running jobs
 - bjobs ⇔ squeue
- Kill jobs
 - bkill ⇔ scancel

SLURM as the batch system for CLAIX 2018

LSF to SLURM transition - general options

	LSF	SLURM
jobname	-J	--job-name (-J)
outputfile	-o	--output (-o)
walltime	-W	--time (-t)
„queue“	-q	--partition (-p)
„project“	-P	--account (-A)

SLURM as the batch system for CLAIX 2018

LSF to SLURM transition - jobtypes

	LSF	SLURM
Number of processes	-n	--ntasks (-n)
shared memory job	-n <num> -a openmp (-R „span[hosts=1]“)	--ntasks=1 --nodes=1 --cpus-per-task=<num>
distributed job	-n <num> -a intelmpi -a openmpi	--ntasks=<num>
hybrid job, n MPI ranks x processes per node with y threads	-n <num> -R „span[ptile=x]“ OMP_NUM_THR...=y	--ntasks=<num> --tasks-per-node=<x> --cpus-per-task=<y>

SLURM as the batch system for CLAIX 2018

CLAIX 2018 partitions

- c18m
 - „mpi“-nodes, Skylake Xeon SP Platinum X8160, 48 cores, 2 sockets, 192 GB RAM
- c18g (future, not yet ready for this workshop)
 - „gpu“-nodes, Skylake Xeon SP Platinum X8160, 48 cores, 2 sockets, 192 GB RAM,
 - 2 Tesla V100, 16GB RAM

Accounts have a default partition and „allowed“ partitions

CLAIX 2018 projects will be able to switch to another partition from the „allowed“ list

Submission without partition parameter results in submission to the projects default partition

SLURM as the batch system for CLAIX 2018

Needed for this workshop

Account „hpclab“ → automatic submission to partition c18m

Two reservations available:

aixcel → general purpose jobs

aixcel_perf → for performance measurements

→ --account=hpclab --reservation=aixcel[_perf]

New defaults for environment modules

Amongst other changes

Per default loaded modules on CLAIX 2018 with SLURM

intel/16.0 → intel/19.0

intelmpi/2017 → intelmpi/2019

openmpi/1.6 → openmpi/3.1

New defaults for environment modules

Amongst other changes

Per default loaded modules on CLAIX 2018 with SLURM

intel/16.0 → intel/19.0	<= API partially changed
intelmpi/2017 → intelmpi/2019	<= did not run on InfiniBand
openmpi/1.6 → openmpi/3.1	<= API changed!

IMPORTANT CHANGE

openmpi/1.10.4 → intelmpi/2019

!!! RECOMPILE YOUR APPLICATIONS FROM SCRATCH !!!

**Thank you very much for your attention.
Questions?**