



aiXcelerate 2017

Paul Kapinos

■ CLAIX is tightly integrated into the RWTH Cluster

- same software, same batch system(, same front ends, ...)
- new CPUs, new fabric
 - old binaries will run
 - **but recompiling/linking all of your software highly recommended**
- you need a project for starting batch jobs on CLAIX
 - you get a test project for the lab session
 - with a project of RWTH or THES class, OR
 - without a project **your jobs will run on old (Bull) hardware**

■ CLAIX front ends:

→ login.hpc.itc.rwth-aachen.de

→ login2.hpc.itc.rwth-aachen.de

→ note the new domain

→ CPU: Intel(R) Xeon(R) CPU E5-2695 v4 @ 2.10GHz

→ AVX, **AVX2**: binaries built here could fail on old (Bull) hardware

■ All other (old) common front ends will do, too!

→ note dedicated front ends e.g. for data transfers

→ older CPUs, no AVX2, binaries built here could be slow on CLAIX

→ Intel compiler: use `-axcode`, e.g. see `$FLAGS_FAST`

- **CLAIX uses the same software stack as (old) Bull cluster**
- **old binaries will run, but could be inefficient**
- **compatibility details:**
 - **Open MPI in versions older than openmpi/1.10.4 (current default) won't utilize Intel OmniPath network at full grade**
 - **(old) binaries optimized for old CPUs (e.g. -xHost, -fast) won't use new CPUs at full grade (no AVX2...)**

- **CLAIX uses the same software stack as Bull cluster**
- **old binaries will run, but could be inefficient**
- **recompiling/linking all of your software highly recommended: use**
 - default compiler (Intel), better: intel/18.0
 - default Open MPI (1.10.4), better: Intel MPI, and
 - -xHost on [login\login2] front end for optimize for CLAIX CPUs
 - won't run on old CPU's,
 - \$FLAGS_FAST (-xSSE4.2 -axCORE-AVX2,AVX)
 - will run fine everywhere (also on old hardware)

- **CLAIX uses the same software stack as Bull cluster**

 - >4 compilers, >2 MPIs, dozens of versions

- **use modules! e.g. for Intel MPI:**

 - `$ module switch openmpi intelmpi`

- **use envvars we set for you! \$FC \$CC \$CXX \$MPIFC \$....**

 - `$ $MPIFC -g $FLAGS_FAST mpihelloworld.f90`

 - `$ $MPIEXEC $FLAGS_MPI_BATCH a.out`

- **consult documentation**

 - <http://www.itc.rwth-aachen.de/primer>

 - <https://doc.itc.rwth-aachen.de/display/CC/>

aiXcelerate 2017

Paul Kapinos | IT Center der RWTH Aachen University

- **CLAIX uses the same software stack as Bull cluster**
- **submit and control you jobs from any front end**
 - you need a project for starting batch jobs on CLAIX
 - without a cleared project your jobs will run on old (Bull) hardware
 - with a cleared project you cannot start jobs on old (Bull) hardware
- **which projects will use CLAIX?**
 - JARA, BUND, NOVA, PREP projects
 - any other project cleared for CLAIX, e.g. for using Pascal GPUs
 - test project for aixCelerate lab session: **hpclab**

```
#!/usr/bin/env zsh
#BSUB -J terminal
#BSUB -W 120
#BSUB -M 1850
#BSUB -S 1300
#BSUB -XF
#BSUB -app vtune
#BSUB -P hpclab
#BSUB -U aixcel_2017

#BSUB -R knl
xterm
```

- <- shebang (!) – you need this for modules!
- <- name of your batch job
- <- time you need (in minutes)
- <- Request memory you need for your job (in MB)
- <- Request stack space (in MB)
- <- Request an (exclusive) job with X11-Forwarding
- <- Request HwC support for VTune/Amplifier
- <- use 'hpclab' project to run on CLAIX
- <- use 'aixCelerate' advanced reservation (optional, on aixCelerate labs only!)
- <- order KNL node type
- <- start an application (a terminal)

- Use Intel MPI
- do not use it on front ends

→ on front ends, MPI ranks are offloaded to back ends (old Hw)

```
#!/usr/bin/env zsh
#BSUB -W 120 -M 1850 -S 1300
#BSUB -app vtune
#BSUB -P hpclab -U aixcel_2017
#BSUB -R knl
#BSUB -a intelmpi
#BSUB -n 4 -R "span[ptile=2]"
module switch openmpi intelmpi; module load intelvtune
$MPIEXEC $FLAGS_MPI_BATCH \
    amplxe-cl -collect hotspots -r result_dir a.out
```



Thank for your attention.

Paul Kapinos <kapinos@itc.rwth-aachen.de>