Quick Start on ATPESC Computing Resources

JaeHyuk Kwack
Argonne National Laboratory

Outline

- The DOE Leadership Computing Facility
- ALCF Aurora/Polaris System
- OLCF Frontier System
- Hands-on

The DOE Leadership Computing Facility

- Collaborative, multi-lab, DOE/SC initiative ranked top national priority in *Facilities for the Future of Science: A Twenty-Year Outlook.*
- Mission: Provide the computational and data science resources required to solve the most important scientific & engineering problems in the world.
- Highly competitive user allocation program (INCITE, ALCC).
- Projects receive 100x more hours than at other generally available centers.
- LCF centers partner with users to enable science & engineering breakthroughs (Liaisons, Catalysts).

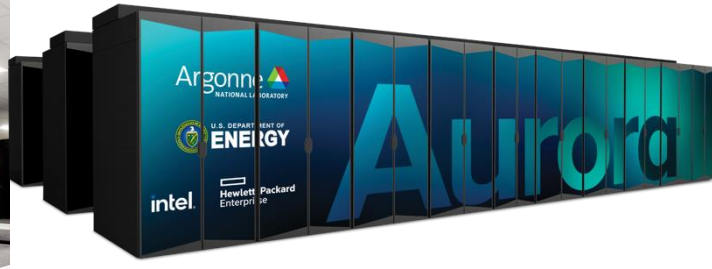

Argonne NATIONAL LABORATORY
OAK RIDGE National Laboratory

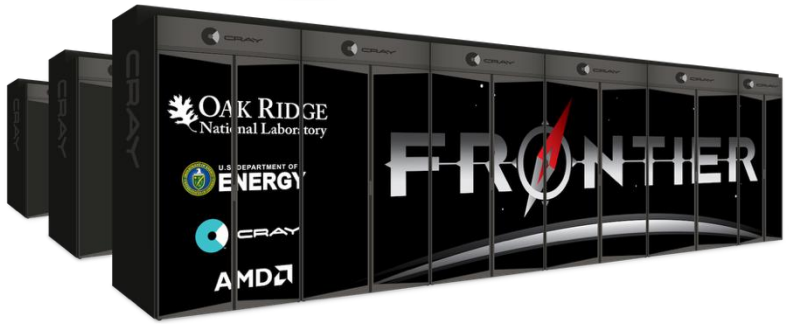

Leadership Computing Facility System

	Argonne LCF		Oak Ridge LCF
System	HPE	HPE	HPE
Name	Polaris	Aurora	Frontier
Compute nodes	560	10,624	9,408
Node architecture	1 x AMD Milan CPU + 4x NVIDIA A100 GPU	2 x Intel Xeon SPR + 6 x Intel PVC GPU	1 x AMD EYPC CPU + 4 x AMD MI250x GPU
Processing Units	560 CPUs + 2,240 GPUs	21,248 CPUs + 63,744 GPUs	9,408 CPUs + 37,362 GPUs
Memory per node, (gigabytes)	512 GB DDR4 + 160 GB HBM2 + 1600 GB SSD	128 GB HBM2e on CPU + 1024 GB DDR5 on CPU + 768 GB HBM2e on GPU	512 GB DDR4 + 512 GB HBM2e
Peak performance	44 Petaflops	> 2 Exaflops DP	> 2 Exaflop DP

AVAILABLE COMPUTING RESOUCES FOR ATPESC

ALCF Systems

- Intel CPUs + Intel PVC GPUs (Aurora)
- AMD CPUs + NVIDIA A100 GPUs (Polaris)

OLCF

- AMD CPUs + AMD MI-250x GPUs (Frontier)

ALCF Aurora/Polaris System

Aurora Exascale Compute Blade

NODE CHARACTERISTICS
6 GPUs - Intel Data Center GPU Max Series
2 CPUs - Intel Xeon CPU Max Series
768 GB GPU HBM Memory
19.66 TB/s Peak GPU HBM BW
128 GB CPU HBM Memory
2.87 TB/s Peak CPU HBM BW
1024 GB CPU DDR5 Memory
0.56 TB/s Peak CPU DDR5 BW
≧ 130 TF Peak Node DP FLOPS
200 GB/s Max Fabric Injection
8 NICs

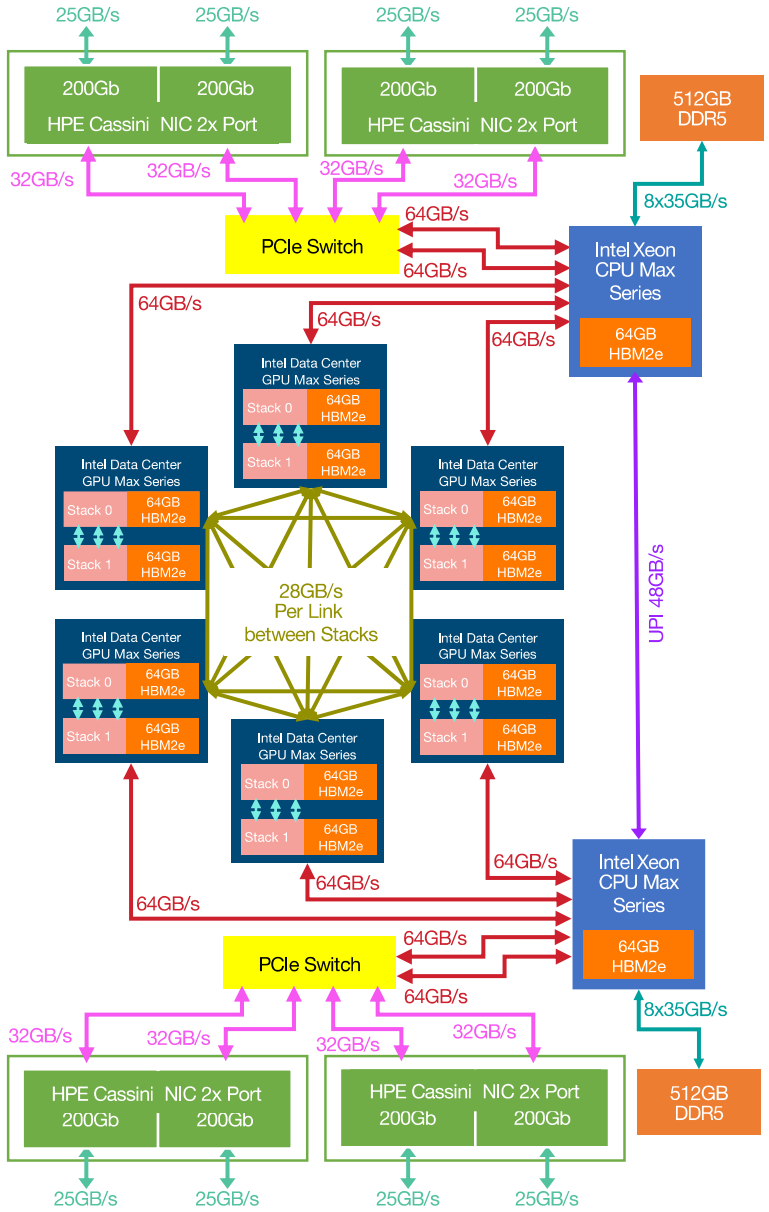

Aurora System Configuration

Intel® Data Center GPU Max 1550 (PVC)

4th Gen Intel XEON Max Series CPU *with High Bandwidth Memory*

Platform
HPE Cray-Ex

Racks - 166
Nodes - 10,624
CPUs - 21,248
GPUs – 63,744

Interconnect
HPE Slingshot 11
Dragonfly topology with adaptive routing
Cassini NIC, 200 Gb/s (25 GB/s), 8 per node
Network Switch:
25.6 Tb/s per switch (64 200 Gb/s ports)
Links with 25 GB/s per direction

Peak FP64 Performance
≧ 2 exaFLOPS

Memory
10.9PiB of DDR @ 5.95 PB/s
1.36PiB of CPU HBM @ 30.5 PB/s
8.16PiB of GPU HBM @ 208.9 PB/s

Network
2.12 PB/s Peak Injection BW
0.69 PB/s Peak Bisection BW

Storage
230PB DAOS Capacity
31 TB/s DAOS Bandwidth

Polaris Single Node Configuration

# of AMD EPYC 7543P CPUs	1
# of NVIDIA A100 GPUs	4
Total HBM2 Memory	160 GB
HBM2 Memory BW per GPU	1.6 TB/s
Total DDR4 Memory	512 GB
DDR4 Memory BW	204.8 GB/s
# OF NVMe SSDs	2
Total NVMe SSD Capacity	3.2 TB
# of Mellanox NICs	2
Total Injection BW (w/ Cassini)	25 (50) GB/s
PCIe Gen4 BW	64 GB/s
NVLink BW	600 GB/s
Total GPU DP Tensor Core Flops	78 TF

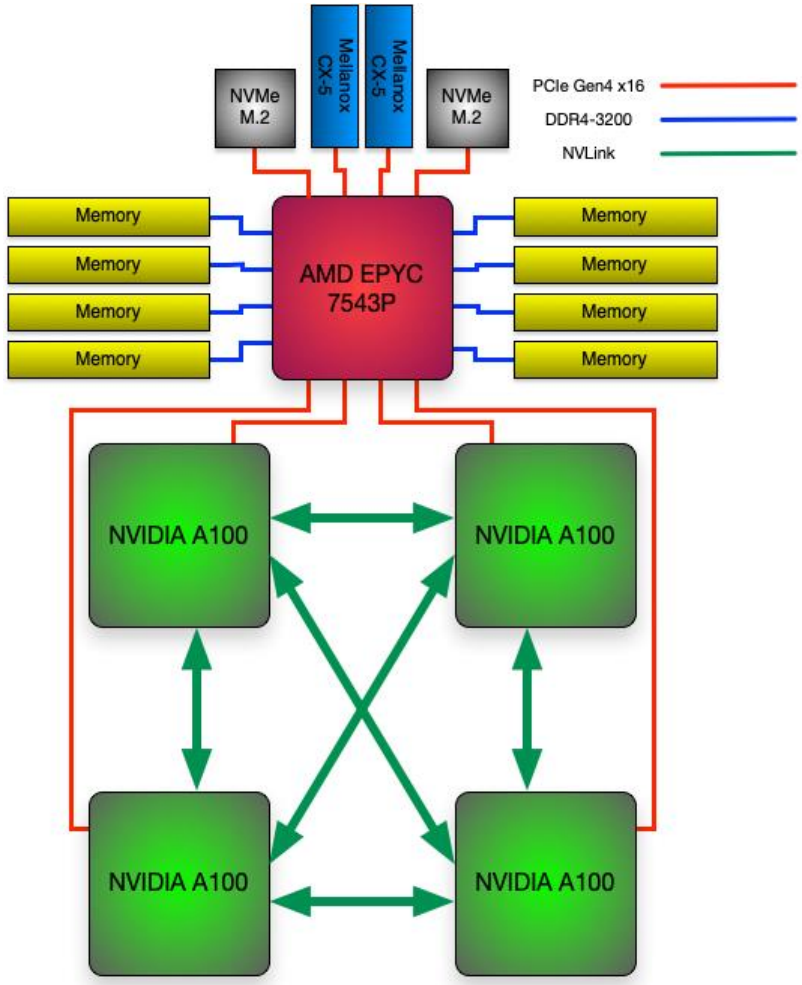

Polaris System Configuration

# of River Compute racks	40
# of Apollo Gen10+ Chassis	280
# of Nodes	560
# of AMD EPYC 7543P CPUs	560
# of NVIDIA A100 GPUs	2240
Total GPU HBM2 Memory	87.5TB
Total CPU DDR4 Memory	280 TB
Total NVMe SSD Capacity	1.75 PB
Interconnect	HPE Slingshot
# of Cassini NICs	1120
# of Rosetta Switches	80
Total Injection BW (w/ Cassini)	28 TB/s 13 TB/s
Total GPU DP Tensor Core Flops	44 PF
Total Power	1.8 MW

Apollo 6500 Gen10+

Aurora/Polaris Filesystems

- Lustre
  - Home directories (/home)
    - Default quota 50GiB
    - Your home directory is backed up
  - Project directory locations
    - Aurora: **/flare/ATPESC2026**
      - **CREATE A SUBDIRECTORY /flare/ATPESC2026/usr/your_username**
    - Polaris: **/eagle/ATPESC2026**
      - **CREATE A SUBDIRECTORY /eagle/ATPESC2026/usr/your_username**
    - Access controlled by unix group of your project
    - Default quota 1TiB
    - Project directories are NOT backed up
  - With large I/O on Lustre, be sure to consider **stripe width**

Aurora/Polaris Modules

- A tool for managing a user's environment
  - Sets your PATH to access desired front-end tools
  - *Your compiler version can be changed here*
- *module commands*
  - *help*
  - *list ← what is currently loaded*
  - *avail*
  - *load*
  - *unload*
  - *switch|swap*
  - *use ← add a directory to MODULEPATH*
  - *display|show*

Aurora Compiling

- Intel oneAPI Environment
  - The oneAPI programming environment is currently the single environment for building and running software to maximally use the available hardware resources. The oneAPI environment is loaded by default for users and is principally defined by the following set of modules and related variants.
    - oneapi: Intel oneAPI HPC toolkit.
    - mpich: MPI libraries.
  - Compiler wrappers
    - mpicc -> icx
    - mpicxx -> icpx
    - mpifort -> ifx
    - Support SYCL and OpenMP target offload
- Libraries found in
  - /opt/aurora/26.26.0/oneapi/
  - /opt/cray

Polaris Compiling

- Cray Programming Environment (PE)
  - HPE provides compiler wrappers by default which includes various libraries (including MPI libraries)
    - Integrates with modules environment
    - HPE provided modules will add headers/libraries/compiler+linker options to compiler
    - -craype-verbose to show actual compile/link command
  - PrgEnv-nvhpc (default)
    - cc -> nvc
    - CC -> nvc++
    - ftn -> nvfortran
    - Support CUDA and OpenMP target offload
    - nvcc still available but not used by wrappers
  - PrgEnv-gnu
    - cc -> gcc
    - CC -> g++
    - ftn -> gfortran
- Libraries found in
  - /opt/nvidia
  - /opt/cray

Aurora/Polaris Running MPI Applications

- Jobs run directly on the compute nodes. The `mpiexec` command runs applications using the Parallel Application Launch Service (PALS)

- mpiexec
  - Execute MPI applications on compute nodes using mpiexec
    - `-n` Total number of MPI ranks
    - `-ppn` Total number of MPI ranks per node
    - `--cpu-bind` CPU binding for application
    - `--depth` Number of CPUs per rank
    - `--env` Set environment variables (e.g., OMP_NUM_THREADS=nthreads)
    - `--hostfile` Indicate file with hostname

- Full list of options available from the man page
  - https://docs.alcf.anl.gov/aurora/running-jobs-aurora/
  - https://docs.alcf.anl.gov/polaris/running-jobs/

Aurora Jobs

- Two parts for running jobs
  - Interacting with scheduler
  - Launching job using `mpiexec`
- Shell script
  - describes parameters for scheduler
  - Commands to run included mpiexec to launch
  - Runs on ‘head’ node of your job
    - Permissible to run computation in your shell script
  - Need to load any of your non-default modules which provide library paths
- `qsub -q prod ./run.sh`
  - Will return the jobid
  - Output and error logs are in submission directory

```
#!/bin/bash
#PBS -A $PROJECT
#PBS -l walltime=01:00:00
#PBS -l select=4
#PBS -l filesystems=home:flare

rpn=12 # assume 1 process per GPU
procs=$((PBS_NODES*rpn))

# job to “run” from your submission directory
cd $PBS_O_WORKDIR

module load <something>

set +x # report all commands to stderr
env
mpiexec -n $procs -ppn $rpn --cpu-bind core -
genvall ./bin <opts>
```

Aurora Interactive job

- Useful for short tests or debugging
- Submit the job with `-I` (letter `I` for Interactive)
  - Debug queue:
    - `qsub -I -l select=1 -l walltime=1:00:00 -q debug -A ATPESC2026 -l filesystems=home:flare`
  - Using reservation:
    - `qsub -I -l select=1 -l walltime=1:00:00 -q ATPESC -A ATPESC2026 -l filesystems=home:flare`
    - Check the reservation queues with `pbs_rstat`
- Wait for job's shell prompt
  - Exit this shell to end your job
- From job's shell prompt, run just like in a script job,
  - `mpiexec -n 4 -ppn 4 ./a.out`
- After job expires, `mpiexec` will fail. *Check* **`qstat $PBS_JOBID`**

Polaris Jobs

- Two parts for running jobs
  - Interacting with scheduler
  - Launching job using `mpiexec`
- Shell script
  - describes parameters for scheduler
  - Commands to run included mpiexec to launch
  - Runs on ‘head’ node of your job
    - Permissible to run computation in your shell script
  - Need to load any of your non-default modules which provide library paths
- `qsub -q prod ./run.sh`
  - Will return the jobid
  - Output and error logs are in submission directory

```
#!/bin/bash
#PBS -A $PROJECT
#PBS -l walltime=01:00:00
#PBS -l select=4
#PBS -l filesystems=home:eagle

rpn=4 # assume 1 process per GPU
procs=$((PBS_NODES*rpn))

# job to “run” from your submission directory
cd $PBS_O_WORKDIR

module load <something>

set +x # report all commands to stderr
env
mpiexec -n $procs -ppn $rpn --cpu-bind core -
genvall ./bin <opts>
```

Polaris Interactive job

- Useful for short tests or debugging
- Submit the job with `-I` (letter `I` for Interactive)
  - Debug queue:
    - `qsub -I -l select=1 -l walltime=1:00:00 -q debug -A ATPESC2026 -l filesystems=home:eagle`
  - Using reservation:
    - `qsub -I -l select=1 -l walltime=1:00:00 -q ATPESC -A ATPESC2026 -l filesystems=home:eagle`
    - Check the reservation queues with `pbs_rstat`
- Wait for job's shell prompt
  - Exit this shell to end your job
- From job's shell prompt, run just like in a script job,
  - `mpiexec -n 4 -ppn 4 ./a.out`
- After job expires, `mpiexec` will fail. *Check* **`qstat $PBS_JOBID`**

Aurora/Polaris Scheduler – PBS Professional

- Primary commands
  - qsub
    - Request resources and start your script on the head node
    - `-A` Allocation
    - `-l` Options
    - `-I` Interactive mode
    - `-q` Which queue to submit otherwise default queue
  - qstat
    - Check on the status of requests
    - `-Q` List queues
    - `-f <jobid>` Detailed information about a job
    - `-x <jobid>` Information about a completed job
  - qalter
    - Update your requests
  - qdel
    - Cancel/delete jobs
  - pbs_rstat
    - Check reservations

Aurora Queues

- Aurora had 5 main queues
  - https://docs.alcf.anl.gov/aurora/running-jobs-aurora/

  - debug
    - 2 nodes max
    - 1 hour max
    - 5 minutes min
  - debug-scaling
    - 2 nodes min
    - 256 nodes max
    - 1 hour max
    - 5 minutes min
  - prod
    - 256 nodes min
    - 10624 nodes max
    - 5 minutes min
    - 24 hours max
  - capacity
    - 1 nodes min
    - 16 nodes max
    - 5 minutes min
    - 7 days max
  - visualization (*by request only*)
    - 1 nodes min
    - 32 nodes max
    - 5 minutes min
    - 8 hours max

Polaris Queues

- Polaris had 3 main queues
  - https://docs.alcf.anl.gov/polaris/running-jobs/
  - debug
    - 2 nodes max
    - 1 hour max
    - 5 minutes min
  - debug-scaling
    - 10 nodes max
    - 1 hour max
    - 5 minutes min
  - prod
    - 10 nodes min
    - 496 nodes max
    - 5 minutes min
    - 24 hours max

Cryptocard tips for ALCF systems

- Tokens
  - Physical Token: The displayed value is a hex string. Type your PIN followed by all letters as CAPITALS.
  - Mobile Token: Type the displayed numbers from your mobile app
- If you fail to authenticate the first time, you may have typed it incorrectly
  - Try again with the **same crypto string** (do NOT refresh button again)
- If you fail again, try a different ALCF host with a fresh crypto #
  - A successful login resets your count of failed logins
- Too many failed logins → your account locked
  - Symptom: You get password prompt but login denied even if it is correct
- Too many failed logins from a given IP → the IP will be blocked
  - Symptom: connection attempt by ssh or web browser will just time out

ALCF References

- Sample files
  - /flare/ATPESC2026/EXAMPLES/track-0-getting-started/ #on Aurora
  - /eagle/ATPESC2026/EXAMPLES/track-0-getting-started/ #on Polaris

- Online docs
  - https://www.alcf.anl.gov/support-center
  - https://docs.alcf.anl.gov/aurora/getting-started-on-aurora/
  - https://docs.alcf.anl.gov/polaris/getting-started/
  - ALCF Beginners Guide (Instructions, examples, and videos)
    - https://github.com/argonne-lcf/ALCFBeginnersGuide/tree/master/aurora
    - https://github.com/argonne-lcf/ALCFBeginnersGuide/tree/master/polaris
    - https://www.alcf.anl.gov/events/getting-started-aurora-and-polaris-bootcamp

OLCF Frontier System

Frontier System overview

- HPE Cray EX Supercomputer architecture
- 74 cabinets, 128 nodes per cabinet (9408 nodes)
- 3rd Gen AMD EPYC 64-core CPU
- 4 AMD Instinct MI250X GPUs
- HPE Slingshot interconnect
- Peak 1.206 Exaflops on HPL benchmark
- Cray, AMD, and GNU software stacks
- Orion - 679 PB multi-tier Lustre filesystem
- NFS storage (/ccs/home, /ccs/proj)

Frontier Compute Node Configuration

- 1x AMD Optimized 3rd Gen EPYC 64 core processor
  - 2 hardware threads per physical core,
  - 2.0GHz base clock, 3.7GHz boost clock
- 512 GB DDR4 memory with 205 GB/s peak bandwidth
- 2x NVMe 2TB SSDs, peak 8 GB/s R, 4 GB/s W, >1.5M IOPs
- 4x AMD MI250X Instinct GPUs
  - 128 GB High-Bandwidth Memory (HBM2E)
  - 3.2 TB/s peak bandwidth
  - 53 TFLOPS double-precision peak for modeling & simulation
  - 2 Graphic Compute Dies (GCDs)
- AMD Infinity Fabric between CPU and GPUs
  - Peak host-to-device (H2D) and device-to-host (D2H) data transfers of 36+36 GB/s per link
- AMD Infinity Fabric between MI250Xs
  - Peak device-to-device bandwidth of 50+50 GB/s per link, low latency
- 4x HPE Slingshot Interconnect 200 GbE NICs
  - Provides 100 GB/s to other nodes, 25 GB/s per port

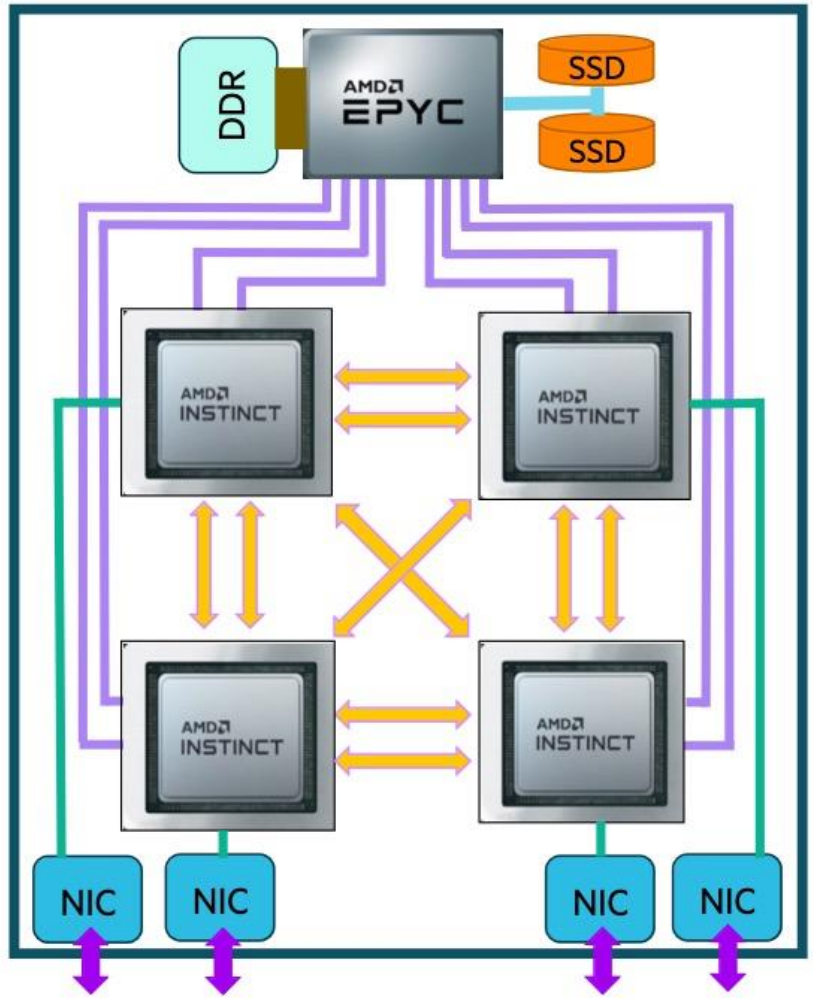

Frontier Module Commands

Command	Description
`module -t list`	Shows a terse list of the currently loaded modules
`module avail`	Shows a table of the currently available modules
`module help <modulename>`	Shows help information about `<modulename>`
`module show <modulename>`	Shows the environment changes made by the `<modulename>` modulefile
`module spider <string>`	Searches all possible modules according to `<string>`
`module load <modulename> [...]`	Loads the given `<modulename>`(s) into the current environment
`module use <path>`	Adds `<path>` to the modulefile search cache and `MODULESPATH`
`module unuse <path>`	Removes `<path>` from the modulefile search cache and `MODULESPATH`
`module purge`	Unloads all modules
`module reset`	Resets loaded modules to system defaults
`module update`	Reloads all currently loaded modules

https://docs.olcf.ornl.gov/systems/frontier_user_guide.html#general-usage

Frontier: Compiling

- Cray, AMD, and GCC compilers are provided through modules. The Cray and AMD compilers are both based on LLVM/Clang. There is also a system/OS versions of GCC available in /usr/bin. The table below lists details about each of the module-provided compilers.

Vendor	Programming Environment	Compiler Module	Language	Compiler Wrapper	Compiler
Cray	`PrgEnv-cray`	`cce`	C	`cc`	`craycc`
			C++	`CC`	`craycxx` or `crayCC`
			Fortran	`ftn`	`crayftn`
AMD	`PrgEnv-amd`	`amd`	C	`cc`	`amdclang`
			C++	`CC`	`amdclang++`
			Fortran	`ftn`	`amdflang`
GCC	`PrgEnv-gnu`	`gcc-native` or `gcc` (<12.3)	C	`cc`	`gcc`
			C++	`CC`	`g++`
			Fortran	`ftn`	`gfortran`

- https://docs.olcf.ornl.gov/systems/frontier_user_guide.html#compiling

Frontier: Slurm Workload Manager

- Slurm Commands (vs. LSF commands)

Command	Action/Task	LSF Equivalent
`squeue`	Show the current queue	`bjobs`
`sbatch`	Submit a batch script	`bsub`
`salloc`	Submit an interactive job	`bsub -Is $SHELL`
`srun`	Launch a parallel job	`jsrun`
`sinfo`	Show node/partition info	`bqueues` or `bhosts`
`sacct`	View accounting information for jobs/job steps	`bacct`
`scancel`	Cancel a job or job step	`bkill`
`scontrol`	View or modify job configuration.	`bstop`, `bresume`, `bmod`

- https://docs.olcf.ornl.gov/systems/frontier_user_guide.html#slurm

Frontier: Batch Scripts

- To submit a batch script,
  - Use the command, `sbatch myjob.sl`
- Description of the example

Line	Description
2	OLCF project to charge
3	Job name
4	Job standard output file (%x: the job name, %j: the Job ID)
5	Walltime requested (in HH:MM:SS format).
6	Partition (queue) to use
7	Number of compute nodes requested
9	Change into the run directory
10	Copy the input file into place
11	Run the job (add layout details)
12	Copy the output file to an appropriate location.

```
1  #!/bin/bash
2  #SBATCH -A ABC123
3  #SBATCH -J RunSim123
4  #SBATCH -o %x-%j.out
5  #SBATCH -t 1:00:00
6  #SBATCH -p batch
7  #SBATCH -N 1024
8
9  cd $MEMBERWORK/abc123/Run.456
10 cp $PROJWORK/abc123/RunData/Input.456 ./Input.456
11 srun ...
12 cp my_output_file $PROJWORK/abc123/RunData/Output.456
```

- https://docs.olcf.ornl.gov/systems/frontier_user_guide.html#batch-scripts

Frontier: Interactive job

- Useful for short tests or debugging
- Submit the job with `salloc`
  - `salloc -A trn001 -J RunSim123 -t 1:00:00 -p batch -N 1`
- Wait for job's shell prompt
  - Exit this shell to end your job
- From job's shell prompt, run just like in a script job,
  - `srun -N 1 -n 8 --ntasks-per-node=8 ./a.out`
- https://docs.olcf.ornl.gov/systems/frontier_user_guide.html#interactive-jobs

Frontier: Monitoring and Modifying Jobs

- Primary commands
  - `scancel`: Cancel or Signal a Job
  - `squeue`: View the queue
    - `squeue -l` : Show all jobs currently in the queue
    - `squeue -l -u $USER` : Show all of your jobs currently in the queue
  - `scontrol show job`: get Detailed Job information
- https://docs.olcf.ornl.gov/systems/frontier_user_guide.html#monitoring-and-modifying-batch-jobs

Frontier: Running MPI Applications

- Jobs run directly on the compute nodes. The `srun` command is used to execute an MPI binary on one or more compute nodes in parallel.
- `srun`
  - `-N` Number of nodes
  - `-n` Total number of MPI tasks
  - `-c` Logical cores per MPI task
  - `--ntasks-per-node=<ntasks>` A maximum count of tasks per node
  - `--gpus` Specify the number of GPUs
  - `--gpu-per-node` Specify the number of GPUs per node

- https://docs.olcf.ornl.gov/systems/frontier_user_guide.html#srun

Cheat Sheets

ATPESC Resources

ALCF – Aurora/ Polaris

-**Project name:** ATPESC2026

-**Note:** use your ALCF Username.

-**Support:** ALCF staff available to help you **via slack**!! and support@alcf.anl.gov

-**Reservations:** Please check the details of the reservations directly on Aurora/Polaris (**command**: *pbs_rstat*)

-**Queue**

  -**check *pbs_rstat*, or default for running without reservation**

-**User guide:**

  -https://www.alcf.anl.gov/support-center

  -https://docs.alcf.anl.gov/aurora/getting-started-on-aurora/

  -https://docs.alcf.anl.gov/polaris/getting-started/

ATPESC Resources

OLCF – Frontier

-**Project name:** trn001

-**Support:** help@olcf.ornl.gov

-**Queue: running without reservation**

-**User guide:**

-**Frontier User Guide:** https://docs.olcf.ornl.gov/systems/frontier_user_guide.html#frontier-user-guide

Questions?

- *Use this presentation as a reference during ATPESC!*
- Supplemental info will be posted as well

Hands-on Exercises

Hands-on exercise

- Aurora hands-on
- Polaris hands-on
- Frontier hands-on

Hands-on exercise: Aurora

```
$ ssh -Y {your_username}@aurora.alcf.anl.gov                    # Login to Aurora
$ module avail                                                  # See available modules
$ module list                                                   # See loaded modules
```

```
jkwack@aurora-uan-0010:~> module list

Currently Loaded Modules:
  1) gcc/13.4.0                3) mpich/opt/5.0.0.aurora_test.3c70a61   5) cray-pals/1.8.0
  2) oneapi/release/2025.3.1   4) libfabric/1.22.0                      6) cray-libpals/1.8.0
```

```
$ qstat -u ${USER}                                              # To see your jobs
$ pbs_rstat                                                     # Check reservation
```

```
jkwack@aurora-uan-0010:~> pbs_rstat
Resv ID         Queue         User     State            Start / Duration / End
--------------------------------------------------------------------------------
M8660168.aurora M8660168      appmm2pb DG     Mon Jul 20 06:00 / 67500 / Tue Jul 21 00:45
M8671302.aurora M8671302      mluczkow DG     Tue Jul 21 11:00 / 604800 / Tue Jul 28 11:00
R8673310.aurora R8673310      mluczkow CO           Fri 10:00 / 86400 / Sat 10:00
R8674622.aurora R8674622      alexku@* CO           Thu 13:00 / 21600 / Thu 19:00
```

Hands-on exercise: Aurora

```
$ qsub -I -l select=1 -l walltime=00:30:00 -l filesystems=home:flare  -A ATPESC2026 -q ATPESC
```

```
jkwack@aurora-uan-0010:~> qsub -I -l select=1 -l walltime=00:30:00 -l filesystems=home:flare  -A ATPESC2026 -q debug
qsub: waiting for job 8674978.aurora-pbs-0001.hostmgmt.cm.aurora.alcf.anl.gov to start
qsub: job 8674978.aurora-pbs-0001.hostmgmt.cm.aurora.alcf.anl.gov ready

jkwack@x4316c3s7b0n0:~>
```

```
$ cd /flare/ATPESC2026/usr/
$ mkdir $USER
$ cd $USER
$ cp -rf /flare/ATPESC2026/EXAMPLES/track-0-getting-started .
$ cd track-0-getting-started/aurora/
$ more Makefile
    …
    CC=mpicc

    hellompi: hellompi.c
          which $(CC)
          $(CC) -g -O0 -o hellompi hellompi.c
    …
```

Hands-on exercise: Aurora

```
$ cat submit.sh
#!/bin/bash
#PBS -l select=1
#PBS -l walltime=00:30:00
##PBS -q debug
#PBS -l filesystems=home:flare
#PBS -A ATPESC2026
#PBS -q ATPESC


cd $PBS_O_WORKDIR


mpiexec -n 4 --ppn 4 ./hellompi
status=$?


echo "mpiexec status is $status"
exit $status
```

Hands-on exercise: Aurora

```
$ mpicc -o hellompi hellompi.c                    # Build the example
$ make clean; make                        # Another way to build the example
```

```
$ mpiexec -n 4 --ppn 4 ./hellompi
```

```
jkwack@x4316c3s7b0n0:/flare/ATPESC2026/usr/jkwack/track-0-getting-started/aurora> mpiexec -n 4 --ppn 4 ./hellompi
3: Hello!
2: Hello!
0: Hello!
1: Hello!
```

```
$ module load xpu-smi
$ xpu-smi discovery
```

More references for Aurora

- https://github.com/argonne-lcf/ALCFBeginnersGuide/tree/master/aurora
- https://www.alcf.anl.gov/events/getting-started-aurora-and-polaris-bootcamp
- https://docs.alcf.anl.gov/aurora/getting-started-on-aurora/

Hands-on exercise: Polaris

```
$ ssh -Y {your_username}@polaris.alcf.anl.gov                    # Login to Polaris
$ module avail                                                   # See available modules
$ module list                                                    # See loaded modules
```

```
jkwack@polaris-login-04:~> module list

Currently Loaded Modules:
  1) nvhpc/23.9          4) cray-mpich/8.1.28   7) cray-libpals/1.3.4  10) libfabric/1.15.2.0      13) darshan/3.4.4
  2) craype/2.7.30       5) cray-pmi/6.1.13     8) craype-x86-milan    11) craype-network-ofi      14) xalt/3.0.2-202408282050
  3) cray-dsmml/0.2.2    6) cray-pals/1.3.4     9) PrgEnv-nvhpc/8.5.0  12) perftools-base/23.12.0
```

```
$ qstat -u ${USER}                                               # To see your jobs
$ pbs_rstat                                                      # Check reservation
```

```
jkwack@polaris-login-04:~> pbs_rstat
Resv ID         Queue         User     State             Start / Duration / End
--------------------------------------------------------------------------------
R5594657.polari R5594657      mluczkow CO     Sun Jul 27 14:30 / 16200 / Sun Jul 27 19:00
R5594658.polari R5594658      mluczkow CO     Mon Jul 28 15:30 / 12600 / Mon Jul 28 19:00
R5594659.polari R5594659      mluczkow CO     Tue Jul 29 08:30 / 37800 / Tue Jul 29 19:00
R5594660.polari R5594660      mluczkow CO     Wed Jul 30 08:30 / 46800 / Wed Jul 30 21:30
R5594661.polari R5594661      mluczkow CO     Thu Jul 31 09:00 / 43200 / Thu Jul 31 21:00
R5594662.polari R5594662      mluczkow CO     Fri Aug 01 08:30 / 36000 / Fri Aug 01 18:30
R5594663.polari R5594663      mluczkow CO     Mon Aug 04 17:30 / 12600 / Mon Aug 04 21:00
R5594664.polari R5594664      mluczkow CO     Tue Aug 05 08:30 / 37800 / Tue Aug 05 19:00
R5594665.polari R5594665      mluczkow CO     Wed Aug 06 08:00 / 39600 / Wed Aug 06 19:00
R5594666.polari R5594666      mluczkow CO     Thu Aug 07 09:00 / 43200 / Thu Aug 07 21:00
R5594667.polari R5594667      mluczkow CO     Wed Jul 30 21:00 / 21600 / Thu Jul 31 03:00
R5594668.polari R5594668      mluczkow CO     Thu Jul 31 21:00 / 21600 / Fri Aug 01 03:00
R5594669.polari R5594669      mluczkow CO     Fri Aug 01 21:00 / 32400 / Sat Aug 02 06:00
R5594670.polari R5594670      mluczkow CO     Mon Aug 04 21:00 / 21600 / Tue Aug 05 03:00
R5594671.polari R5594671      mluczkow CO     Tue Aug 05 21:00 / 21600 / Wed Aug 06 03:00
R5594672.polari R5594672      mluczkow CO     Wed Aug 06 21:30 / 30600 / Thu Aug 07 06:00
```

Hands-on exercise: Polaris

```
$ qsub -I -l select=1 -l walltime=00:30:00 -l filesystems=home:eagle  -A ATPESC2026 -q ATPESC
```

```
jkwack@polaris-login-04:~> qsub -I -l select=1 -l walltime=00:30:00 -l filesystems=home:eagle  -A ATPESC2026 -q debug
qsub: waiting for job 7257917.polaris-pbs-01.hsn.cm.polaris.alcf.anl.gov to start
qsub: job 7257917.polaris-pbs-01.hsn.cm.polaris.alcf.anl.gov ready


Currently Loaded Modules:
  1) nvidia/25.5       4) cray-mpich/9.0.1   7) cray-libpals/1.7.1  10) xalt/3.2.1-1-g67d304d4-202511121555  13) craype-network-ofi
  2) craype/2.7.35     5) cray-pmi/6.1.16    8) craype-x86-milan    11) PrgEnv-nvidia/8.6.0                 14) perftools-base/25.09.0
  3) cray-dsmml/0.3.1  6) cray-pals/1.7.1    9) darshan/3.4.4       12) libfabric/2.2.0rc1
```

```
$ cd /eagle/ATPESC2026/usr/
$ mkdir $USER
$ cd $USER
$ cp -rf /eagle/ATPESC2026/EXAMPLES/track-0-getting-started .
$ cd track-0-getting-started/polaris/
$ more Makefile
    …
    CC=cc

    hellompi: hellompi.c
          which $(CC)
          $(CC) -g -O0 -o hellompi hellompi.c
    …
```

Hands-on exercise: Polaris

```
$ cat submit.sh
#!/bin/bash
#PBS -l select=1
#PBS -l walltime=00:30:00
##PBS -q debug
#PBS -l filesystems=home:eagle
#PBS -A ATPESC2026
#PBS -q ATPESC


cd $PBS_O_WORKDIR


mpiexec -n 4 --ppn 4 ./hellompi
status=$?


echo "mpiexec status is $status"
exit $status
```

Hands-on exercise: Polaris

```
$ cc -o hellompi hellompi.c             # Build the example
$ make clean; make                      # Another way to build the example
```

```
$ mpiexec -n 4 --ppn 4 ./hellompi
```

```
jkwack@x3004c0s31b0n0:/eagle/ATPESC2026/usr/jkwack/track-0-getting-started/polaris> mpiexec -n 4 --ppn 4 ./hellompi
0: Hello!
1: Hello!
2: Hello!
3: Hello!
```

```
$ nvidia-smi
```

More references for Polaris

- https://github.com/argonne-lcf/ALCFBeginnersGuide/tree/master/polaris
- https://www.alcf.anl.gov/events/getting-started-aurora-and-polaris-bootcamp
- https://docs.alcf.anl.gov/polaris/getting-started/

Hands-on exercise: Frontier

```
$ ssh -Y {your_username}@frontier.olcf.ornl.gov                    # Login to Frontier
$ module avail                                                     # See available modules
$ module list                                                      # See loaded modules
```

```
[jkwack@login07.frontier ~]$ module list

Currently Loaded Modules:
  1) craype-x86-trento                2) libfabric/2.3.1                  3) craype-network-ofi
  4) perftools-base/24.11.0           5) xpmem/1.0.1-1.5_1_gfb6998056825  6) cray-pmi/6.1.15
  7) cce/18.0.1                       8) craype/2.7.33                    9) cray-dsmml/0.3.0
 10) cray-mpich/8.1.31               11) cray-libsci/24.11.0             12) PrgEnv-cray/8.6.0
 13) Core/25.03                      14) tmux/3.4                        15) darshan-runtime/3.4.6-mpi (E4S)
 16) hsi/default                     17) lfs-wrapper/0.0.1               18) DefApps
 19) craype-accel-amd-gfx90a         20) rocm/6.2.4

  Where:
   E4S:  E4S: Extreme-scale Scientific Software Stack (E4S) https://e4s.io/index.html
```

```
$ squeue -l -u $USER                                               # To see your jobs
```

Hands-on exercise: Frontier

```
$ salloc -A trn001 -t 1:00:00 -p batch -N 1
```

```
[jkwack@login03.frontier ~]$ salloc -A trn001 -t 1:00:00 -p batch -N 1
salloc: Pending job allocation 5070510
salloc: job 5070510 queued and waiting for resources
```

```
$ cp -r /ccs/proj/trn001/track-0-getting-started .
$ cd track-0-getting-started/frontier/
$ more Makefile
    …
    CC=cc

    hellompi: hellompi.c
          which $(CC)
          $(CC) -g -O0 -o hellompi hellompi.c
    …
```

Hands-on exercise: Frontier

```
$ more submit.sh
    #!/bin/bash
    #SBATCH -p batch
    #SBATCH -A trn001
    #SBATCH -N 1
    #SBATCH -t 60:00

    srun -n 8  ./hellompi
    status=$?

    echo "mpiexec status is $status"
    exit $status
$ cc -o hellompi hellompi.c             # Build the example
$ make clean; make                      # Another way to build the example
$ srun -n 4 ./hellompi
```

- More references for Frontier
  - https://docs.olcf.ornl.gov/systems/frontier_user_guide.html

Thank you!

Supplemental Info

-

ARGONNE TRAINING PROGRAM ON EXTREME-SCALE COMPUTING

Produced by Argonne National Laboratory, a U.S. Department of Energy Laboratory managed by UChicagoArgonne, LLC under contract DE-AC02-06CH11357.

Special thanks to the National Energy Research Scientific Computing Center (NERSC) and Oak Ridge Leadership Computing Facility (OLCF) for the use of their resources during the training event.