جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

SHAHEEN SUPERCOMPUTING LABORATORY

Shaheen II Get Started

Shaheen II Spec:

System	**Cray XC40 with 36 cabinets**
Processor type	Intel Haswell 2.3GHz, 2 CPU sockets per node, 16 processors cores per CPU
Total Nodes	6174 nodes
Total Cores	197,568 cores
Memory	128 GB of memory per node, over 790 TB total memory
Interconnect	Cray Aries with Dragonfly topology
Scheduler	SLURM
Storage	Lustre parallel file system with 17.4 PB
Burst Buffer	1.5 PB
Linux Release	Cray Linux Environment CLE 7

To login:

```
$ ssh <username>@shaheen.hpc.kaust.edu.sa
```

Login with "-X" or "-Y" to enable X11 forwarding.

To compile:

Three programming environments are supported. PrgEnv-cray (default), PrgEnv-intel, and PrgEnv-gnu. Use module swap to change PrgEnv, e.g.

```
$ module swap PrgEnv-cray PrgEnv-intel
```

Use the compiler driver wrappers cc, CC, ftn to compile and link C, C++, and Fortran codes, respectively. The wrappers are the same for all programming environments. For example

C	cc –c <any_other_flags> prog.c
C++	CC –c <any_other_flags> prog.cpp
Fortran	ftn –c <any_other_flags> prog.f90

Within a programming environment a user can switch between different compiler versions.

```
$ module swap gcc gcc/8.1.0
```

Scheduler and Queues

To run:

- SLURM is the batch scheduler. The following is a basic example of a batch script:

```
#!/bin/bash
#SBATCH -account=kxxxx
#SBATCH -job-name=job_name
#SBATCH -output=job_name.out
#SBATCH -error=job_name.err
#SBATCH -nodes=4
#SBATCH -time=00:30:00
srun --ntasks=128 --hint=nomultithread
--ntasks-per-node=32  --ntasks-per-socket=16 ./exe
```

- Hyperthreading is enabled by default
- Generator of jobs scripts:
  https://www.hpc.kaust.edu.sa/job
- Launch jobs with **`sbatch myjobscript.sh`**
- Cancel job : **`scancel job_id`**

Queues:

- Use "**sinfo**" for the queue status and "**squeue**" to observe your job status.
- **workq**: Default queue with a maximum 24 hours wall clock time.
- **debug**: Default queue with a maximum of 4 nodes and 30 minutes wall clock time. Submit with **`sbatch --partition=debug`** in the command line or add **`#SBATCH -partition=debug`** in the job script
- **72hours**: Queue with a maximum of 72 hours wall clock time. Add the following in your job script to use the 72 hours queue. This queue is subject to approval.

```
#SBATCH -partition=72hours
#SBATCH -qos=72hours
```

Storage, Quotas, Allocations

To store:

- Compute nodes can access only **/scratch** and **/project** directories. Jobs submitted from /home will fail.
- **/home/<username>**: Home directory, designed for development, quota of 200GB. Previous versions of files can be recovered from /home/<username>/snapshot directory.
- **/scratch/<username>**: Temporary individual storage for data needed for execution. Files not accessed in the last 60 days will be deleted.
- **/project/kxxxx**: Project directory **for** medium term. Each PI has a default limit of 80TB.
- **/scratch/tmp**: temporary folder that will be cleaned every 3 days.
- User are limited to 1M files on /project and /scratch

To check:

- Your Group information, use the "**groupies**" command
- Your quota on Lustre: **`lfs quota -uh $USER /lustre`**
- Your allocation information use the "**sb**" command, e.g.

```
sb k1xxx
Project kxx: Title of Project
PI: Name

Allocations      Core hours
---------------------------
2020-02-25           250000
---------------------------
Expired on       2021-02-25
---------------------------
Allocated            250000
Shaheen              229939
Neser                     0
---------------------------
Balance                   0
---------------------------
```

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

SHAHEEN
SUPERCOMPUTING LABORATORY

Compiler Flags

Feature	Cray	Intel	GNU
Recommended compiler optimization level	default (-O3)	default (-O2)	-O3 -ffast-math
Aggressive Optimization	-O3 -hfp3	-Ofast -fp-model fast=2	-O3 -ffast-math
Activate OpenMP directives and pragmas in the code	-homp	-qoenmp	-fopenmp
Desactivate OpenMP	-hnoomp		
Read and write Fortran unformatted data files as big-endian	-h byteswapio	-convert big_endian	-fconvert=swap
Process Fortran source using fixed form specifications.	-f fixed	-fixed	-ffixed-form
Process Fortran source using free form specifications.	-f free	-free	-ffree-form
Show version number of the compiler.	-V	--version	--version
Zero fill all uninitialized variables.	-h zero	not implemented	-finit-local-zero
Creates .mod files to hold Fortran90 module information for future compiles	-e m		
Specifies the directory to which *file*.mod files are written when the -e m option is specified	-J *dir_name*		
Listing compiler feedbacks , produces .lst files	-hlist=a	-opt-report3	-fdump-tree-all

- CCE is based on LLVM. For classical CCE, use cce-classic
- More information on the CLE7 and environment: https://www.hpc.kaust.edu.sa/CLE7_SLES15

Software & Libraries

Before requesting the installation of new packages or libraries, please check if the desired package is already installed on the system.

- To find the list of all the packages installed:
  ```
  $ module avail
  ```
- To find a specific package:
  ```
  $ module avail -S name
  ```
- To get information on the package usage:
  ```
  $ module help xxxx
  $ module show xxxx
  ```
- To display Cray Scientific Libraries:
  ```
  $ module avail -L
  ```
- To load a module: module:
  ```
  $ module load xxxxx
  ```

Here is a selection of libraries and applications already installed on Shaheen II:

- I/O Libraries
  - HDF5, NetCDF
- Numerical Libraries
  - LIBSCI, PETCS, FFTW, MKL, …
- Visualization Tools
  - Gnuplot, Paraview
- Debugging Tools
  - lgdb, atp, ARM(DDT), Totalview, STAT
- Performance tools
  - Craypat, ARM(DDT), PAPI
- Some Third party Software
  - VASP, CP2K, NAMD, LAMMPS, …
- List of application: https://www.hpc.kaust.edu.sa/app7

General Tips

- Currently, dynamic linking is the default. To switch between different link types you can set CRAYPE_LINK_TYPE to "static" or pass the "-static" or "-dynamic" option to the linking wrapper (cc, CC or ftn).
- LIBSCI is the collection of numerical routines optimized for best performance on Cray systems. It gathers BLAS, LAPACK, SCALAPACK and is highly recommended to be used instead of your own versions.
- When calling libraries installed by Cray, such as LIBSCI, HDF5, NetCDF you do not need to add –l, -L and –l flags. Instead, you will have to remove these paths from your Makefiles.
- Default I/O striping is 1, optimal for many cases especially when every MPI process writes to its own file resulting in as many files as number of processes used.
- Increase the stripe count when multiple processes write to a single shared file as with MPI-IO and HDF5 or NetCDF. Use the following command with a maximum stripe count of 144:

```
$ lfs setstripe  -c [stripe-count] filename/directory
```

To get an account:

KAUST members should fill-in the Individual Access Application (IAA) and the Project Proposal(PP) forms. Forms are available at::

https://www.hpc.kaust.edu.sa/account-applications

For more information:

- Please visit the user guide and training materials at: http://hpc.kaust.edu.sa/
- Please email for any question and issues to help@hpc.kaust.edu.sa.
- Follow us on Twitter: https://twitter.com/KAUST_HPC

@April 2020