

# Getting Started on ThetaGPU

**Colleen Bertoni**  
**Argonne Leadership Computing Facility**

# Outline


<https://www.alcf.anl.gov/user-guides>

- ThetaGPU (DGX A100)
  - System Overview
  - Software & Environment Modules
  - Building your code
  - Queuing and running jobs with qsub




# ThetaGPU

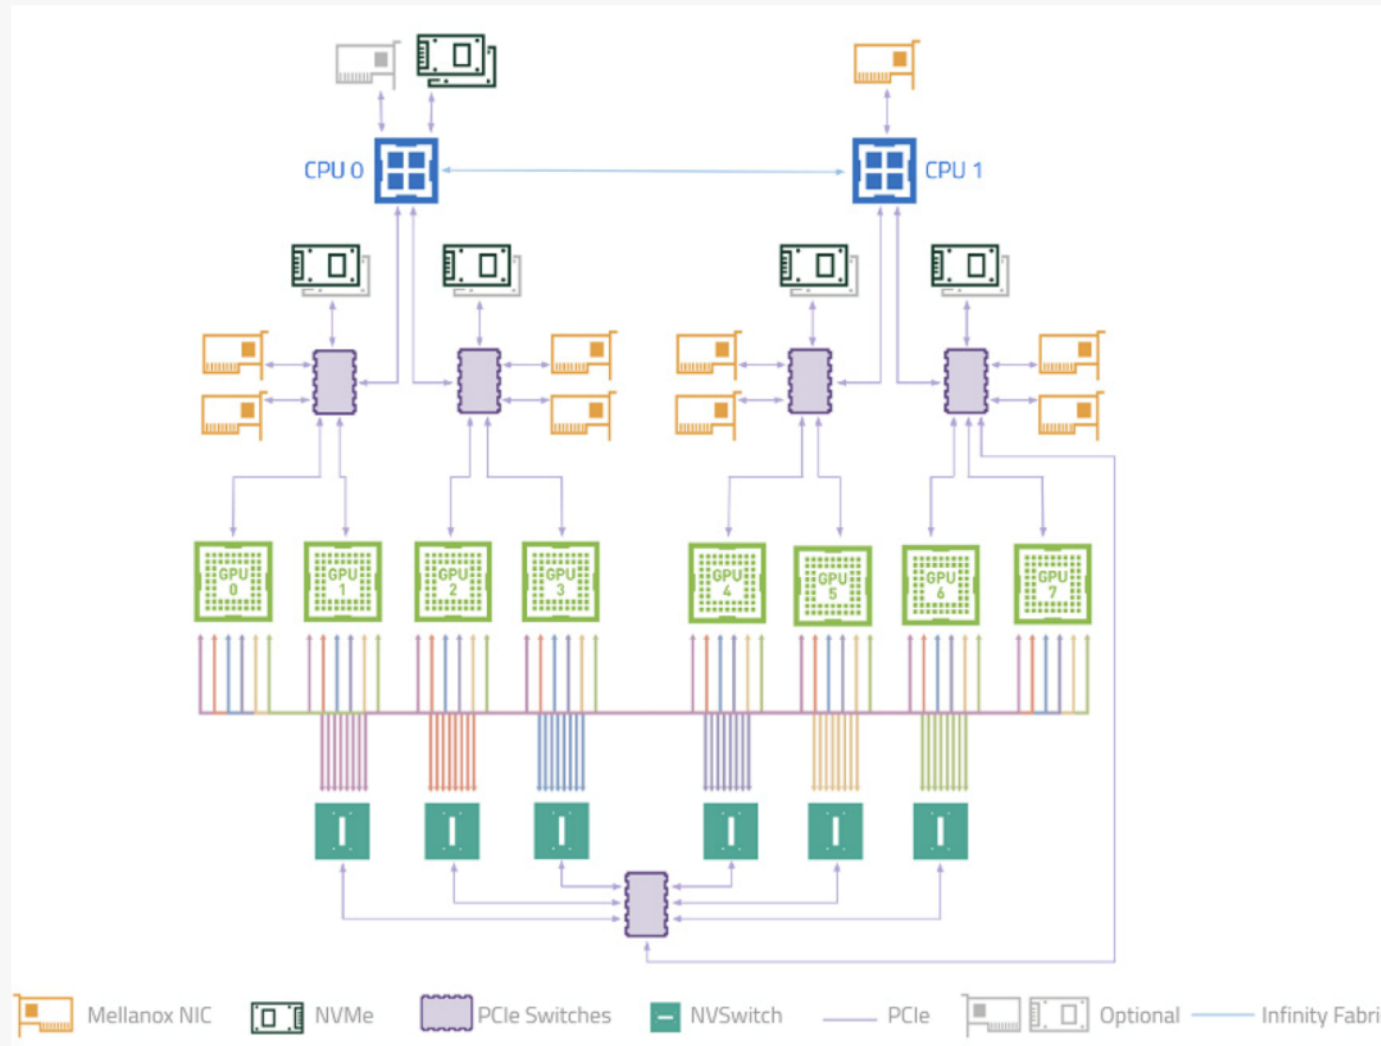
<https://www.alcf.anl.gov/theta>



- Theta expansion to support coronavirus research that is now open for general use
- NVIDIA DGX A100 partition
  - 24 nodes each with
    - 8 NVIDIA A100 Tensor Core GPUs & 320 GB HBM memory
    - 2 AMD Rome 64-core CPUs & 1 TB DDR4
    - 15 TB SSD (4 x 3.84 TB), 25 Gb/s bandwidth
    - 8 HDR 200 NICs (compute network)
    - 2 HDR 200 NICs (storage network)
- Dedicated Compute Fabric
  - 20 Mellanox QM9700 HDR200 40-port switches in fat-tree topology
- Project filesystem is Theta's 10 PB Lustre with 210 GB/s throughput



# ThetaGPU – Node Overview



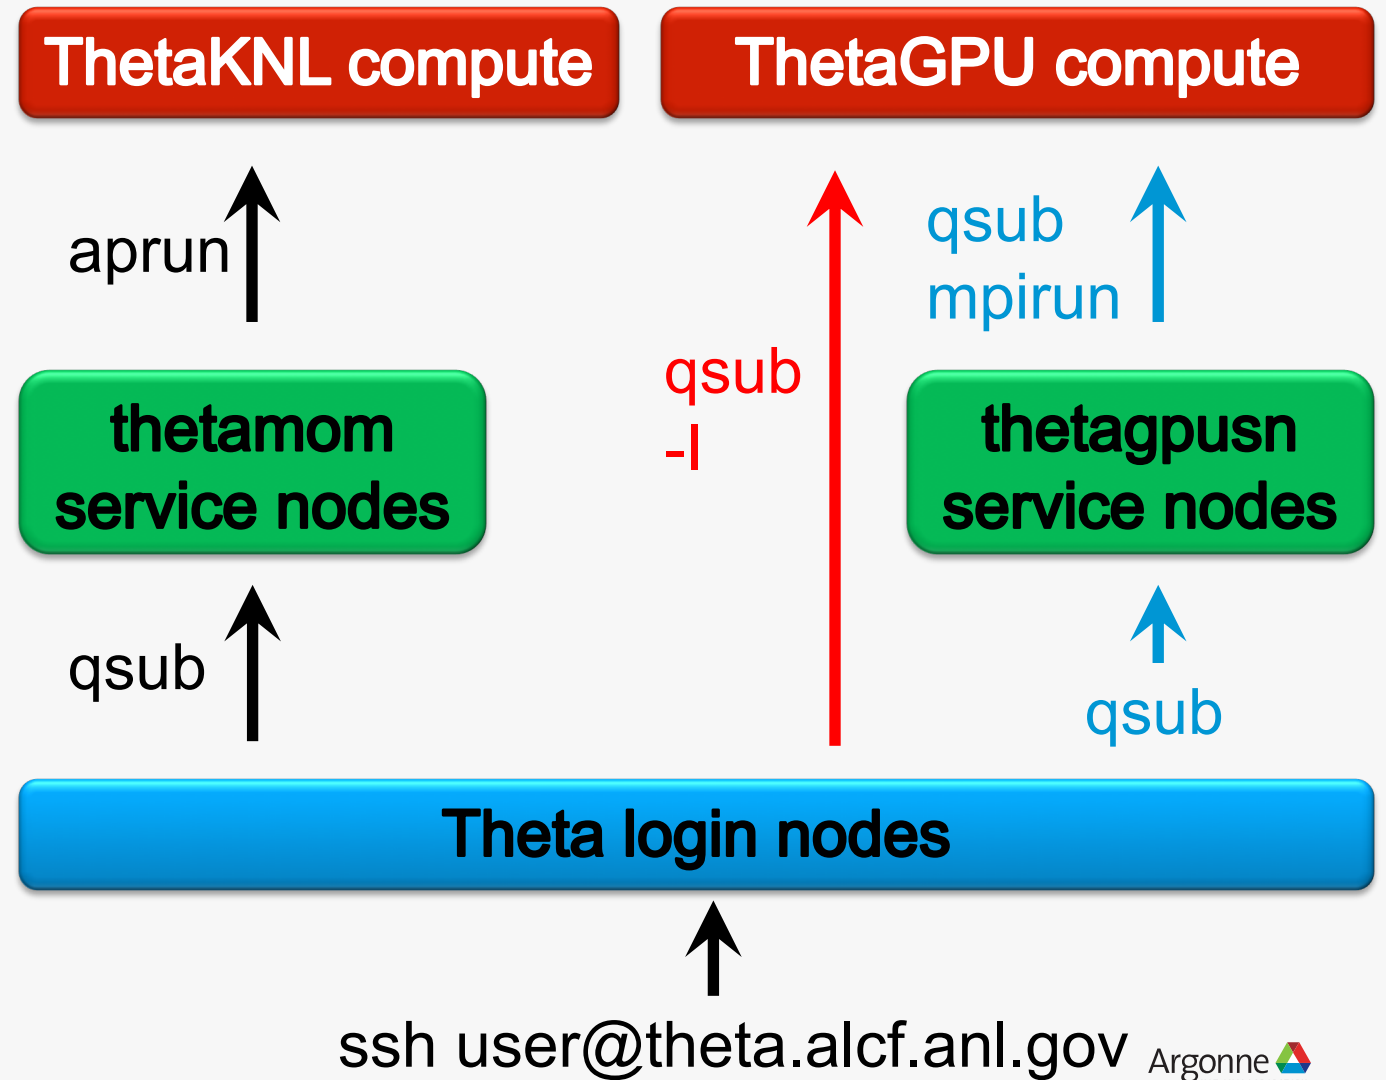
<https://docs.nvidia.com/dgx/pdf/dgxa100-user-guide.pdf>

<https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/dgx-a100/dgxa100-system-architecture-white-paper.pdf>

# ThetaGPU - Logging in and Environment

<https://www.alcf.anl.gov/support-center/theta/theta-thetagpu-overview#theta-gpu>

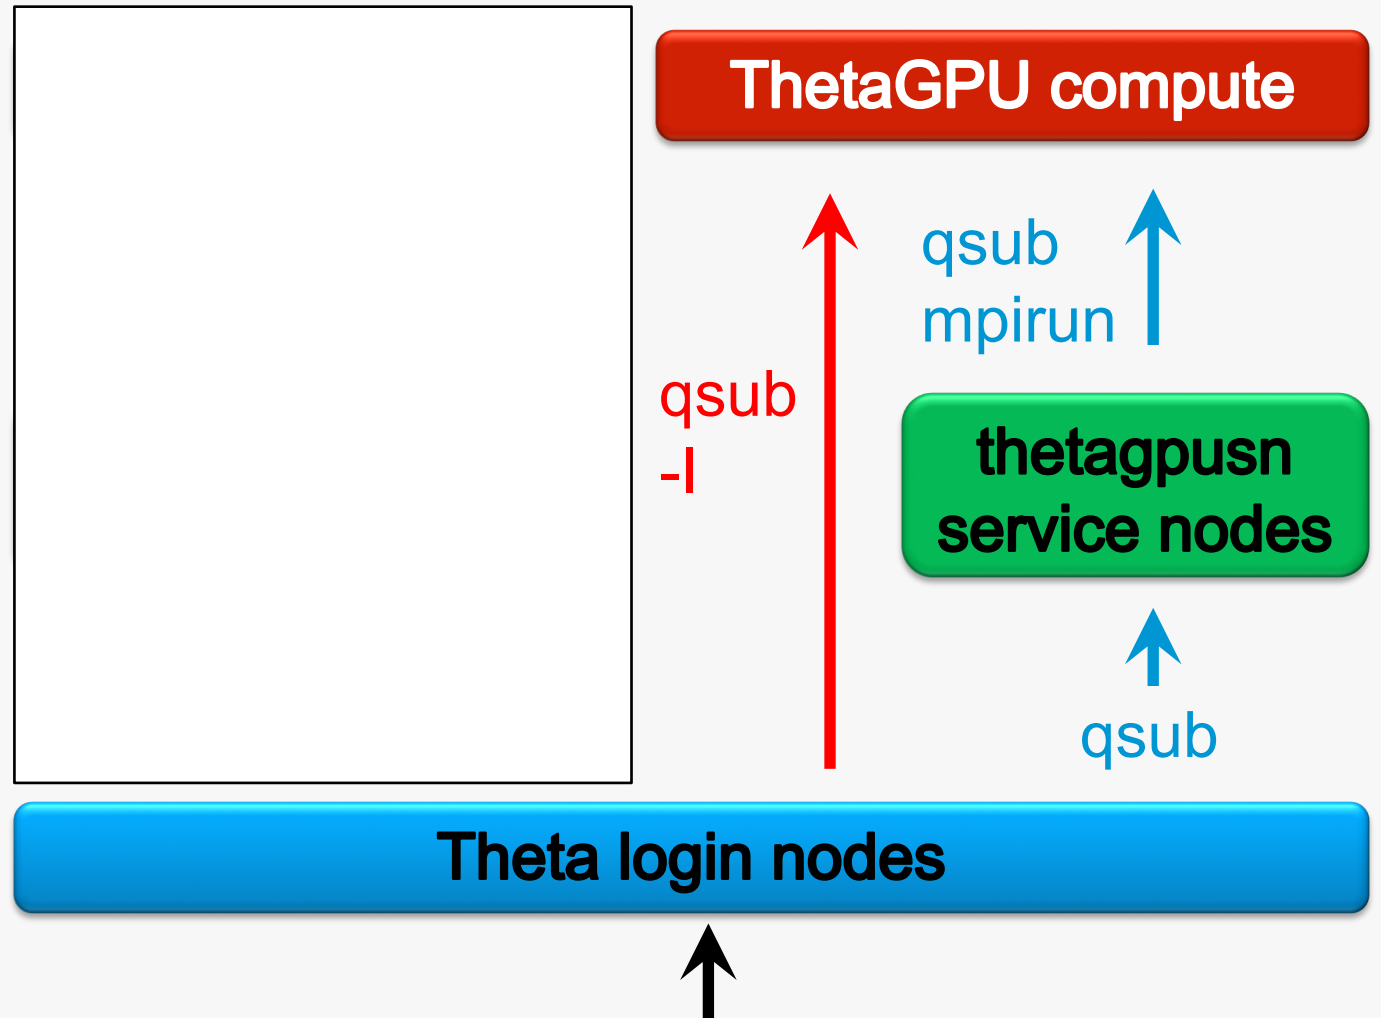
- Use Theta login nodes  
`$ ssh user@theta.alcf.anl.gov`
- Load ThetaGPU scheduler  
`$ module load cobalt/cobalt-gpu`
- Use ThetaGPU compute nodes for building and development  
`$ qsub -l -n 1 -t 60 -q full-node -A ...`
- Can also login to ThetaGPU service nodes, if needed  
`$ ssh thetagpusn1`  
`$ qsub -l -n 1 -t 60 -q full-node -A ...`



# ThetaGPU - Logging in and Environment

<https://www.alcf.anl.gov/support-center/theta/theta-thetagpu-overview#theta-gpu>

- Use Theta login nodes  
`$ ssh user@theta.alcf.anl.gov`
- Load ThetaGPU scheduler  
`$ module load cobalt/cobalt-gpu`
- Use ThetaGPU compute nodes for building and development  
`$ qsub -l -n 1 -t 60 -q full-node -A ...`
- Can also login to ThetaGPU service nodes, if needed  
`$ ssh thetagpusn1`  
`$ qsub -l -n 1 -t 60 -q full-node -A ...`



ssh user@theta.alcf.anl.gov

# ThetaGPU - Software & Libraries

We recommend Nvidia or LLVM compilers, GNU compilers do not currently offload

<https://www.alcf.anl.gov/support-center/theta-gpu-nodes>

- List available modules on ThetaGPU compute node

thetagpu##\$ module avail

```
----- /usr/local/lmod/lmod/modulefiles -----
Core/lmod   Core/settarg

----- /lus/theta-fs0/software/environment/thetagpu/lmod/modulefiles -----
Core/StdEnv      (L,D)  conda/tensorflow/2020-12-23      llvm/release-12.0.0      (D)  nvhpc-nompi/21.3
aocl/blis-3.0    conda/tensorflow/2021-01-08      nccl/nccl-v2.8.4-1_CUDA11      nvhpc/20.9      (D)
cmake/3.19.5     conda/tensorflow/2021-03-02 (D)  nvhpc-byo-compiler/20.9      (D)  nvhpc/21.2
conda/pytorch/2020-11-25  hdf5/1.8.13      nvhpc-byo-compiler/21.2      nvhpc/21.3
conda/pytorch/2021-03-02 (D)  hdf5/1.12.0      (D)  nvhpc-byo-compiler/21.3      openmpi/openmpi-4.0.5      (L)
conda/tensorflow/2020-11-11  llvm/main-20210112      nvhpc-nompi/20.9      (D)  openmpi/openmpi-4.1.0_ucx-1.10.0
conda/tensorflow/2020-12-17  llvm/main-20210426      nvhpc-nompi/21.2      (D)  openmpi/openmpi-4.1.0      (D)

----- /lus/theta-fs0/software/spack/share/spack/modules/linux-ubuntu18.04-x86_64 -----
autoconf-2.69-gcc-7.5.0-wmttzuv      gdbm-1.18.1-gcc-10.2.0-ia4egqb      mpfr-4.0.2-gcc-7.5.0-mpv2v7v      readline-8.0-gcc-10.2.0-ephdh34
autoconf-archive-2019.01.06-gcc-7.5.0-bdyarrk      gdbm-1.18.1-gcc-7.5.0-4av4gyw      ncurses-6.2-gcc-10.2.0-qjpgcs6      readline-8.0-gcc-7.5.0-t54jzdy
automake-1.16.3-gcc-7.5.0-stmktof      gmp-6.1.2-gcc-7.5.0-3ol3tld      ncurses-6.2-gcc-7.5.0-crhlefo      zlib-1.2.11-gcc-10.2.0-glt2u7u
berkeley-db-18.1.40-gcc-10.2.0-fle5h4p      libiconv-1.16-gcc-7.5.0-jearpk4      openssl-1.1.1j-gcc-10.2.0-3g4hwmwz      zlib-1.2.11-gcc-7.5.0-smoyzzo
berkeley-db-18.1.40-gcc-7.5.0-vd7vwr5      libsigsegv-2.12-gcc-7.5.0-lbrx7ln      perl-5.32.1-gcc-10.2.0-grji3ix      zstd-1.4.5-gcc-7.5.0-rnf7xyj
cmake-3.19.5-gcc-10.2.0-felctqr      libtool-2.4.6-gcc-7.5.0-jdxbjft      perl-5.32.1-gcc-7.5.0-op6xocu
diffutils-3.7-gcc-7.5.0-otkkten      m4-1.4.18-gcc-7.5.0-mkc3u4x      pkgconf-1.7.3-gcc-10.2.0-4aysapw
gcc-10.2.0-gcc-7.5.0-jj2fh4j      mpc-1.1.0-gcc-7.5.0-pj4yncj      pkgconf-1.7.3-gcc-7.5.0-4sh6pym

Where:
L: Module is loaded
D: Default Module

Use "module spider" to find all possible modules and extensions.
Use "module keyword key1 key2 ..." to search for all possible modules matching any of the "keys".
```

# ThetaGPU - NVIDIA Compilers

<https://www.alcf.anl.gov/support-center/theta-gpu-nodes/compiling-and-linking-thetagpu>

- NVIDIA HPC SDK
  - Load module, single node
    - \$ module load nvhpc
    - \$ module list
    - Currently Loaded Modules:
      - 1) Core/StdEnv 2) nvhpc/21.3 3) openmpi-4.1.0\_nvhpc-21.3
  - GPU Programming Models: CUDA, OpenMP
  - Use nvc++, nvc, nvfortran
  - For OpenMP offload: -mp=gpu -gpu=cc80
  - Use mpicxx, mpicxx, mpif90, etc...

<https://developer.nvidia.com/hpc-sdk>



# ThetaGPU - LLVM Compilers

<https://www.alcf.anl.gov/support-center/theta-gpu-nodes/compiling-and-linking-thetagpu>

- LLVM w/ OpenMP offload

- Load module

```
$ module load llvm
```

```
$ module list
```

Currently Loaded Modules:

1) openmpi/openmpi-4.0.5 2) Core/StdEnv 3) llvm/release-12.0.0

- GPU Programming Models: CUDA, OpenMP
- Use clang++, clang
- For OpenMP offload: -fopenmp -fopenmp-targets=nvptx64
- Use MPI wrappers: mpicxx, mpicc, etc...

# ThetaGPU – Running a Job

<https://www.alcf.anl.gov/user-guides/cobalt-job-control-xc40>

```
ssh user@theta.alcf.anl.gov
module load cobalt/cobalt-gpu
# submit to the reservation, using two possible queues:
# ecp_openmp_hackathon queue (full node, 8 GPUs):
# limits are currently set at 1hr max walltime, 1 job max running, and 5 jobs max queued
qsub -I -q ecp_openmp_hackathon -t 60 -n 1 -A ecp_omp2021
```

```
# single-gpu queue (1 GPU):
# limits are currently set at 1hr max walltime, 1 job max running, and 1 jobs max queued
qsub -I -q single-gpu -t 60 -n 1 -A ecp_omp2021
```

# Examples for OpenMP offload are here:

[https://github.com/argonne-lcf/CompPerfWorkshop-2021/tree/main/01\\_openmp/demo](https://github.com/argonne-lcf/CompPerfWorkshop-2021/tree/main/01_openmp/demo)



# ANY QUESTIONS?