



Pelin Blanton

Wake Forest University - Senior
Computer Science, Mathematics

Logan Jacobs

Wake Forest University - Senior
Computer Science

Michael Wang

Wake Forest University - Senior
Computer Science, Applied Statistics

Ruby Flanagan

Wake Forest University - Senior
Computer Science, Mathematics

Christopher Lawrence

Wake Forest University - Senior
Biology

Xena Eriksen

Salem College - Senior
Health Sciences, Mathematics

Preparation

- All team members have experience with WFU's supercomputing cluster, DEAC, through academic coursework and research fellowships
- 4 out of 6 team members are enrolled in a new Architecting HPC Clusters course meeting twice a week
- Using Ansible playbooks and roles to configure a Slurm cluster on Jetstream2's Rocky Linux 9 image
- Using architectures across Jetstream2, WFU's DEAC Cluster, and Amazon Web Services to experiment with different software build scripts and optimizations

Team Intro

- We are from Winston-Salem, North Carolina representing both Wake Forest University (WFU) and Salem Academy and College
- The team name "Daemon Deacons" is a play on WFU's mascot, the "Demon Deacon"
- All current team members are SCC first-timers
- Equal representation of men and women
- Variety of majors and minors with emphasis on Computer Science and Mathematics

Software

- Rocky Linux 9 OS
- GNU GCC Compiler
- AOCL Libraries
- Open MPI

To optimize our software on Jetstream2 we're utilizing the GNU GCC compiler and Open MPI to achieve a base level of performance while utilizing the AOCL libraries to take advantage of the underlying AMD CPUs available on our CPU and A100 GPU compute nodes.

Strategies

- Provision Jetstream2 resources to maximize quotas for CPU and RAM

Login
m3.quad

GPU
g3.xl

Compute
m3.2xl

m3.2xl

m3.2xl

r3.large

- Primary cluster configuration being composed of one login node, one GPU node, and three CPU compute nodes
- Flexibility to shelve two m3.2xl instances for one r3.large instance for high memory applications
- Utilize Slurm partitions to direct jobs to individual compute nodes
- Scale CPU applications with Open MPI, but primarily focus on single node jobs to reduce latency