

# SC19 Student Cluster Challenge Final Architecture Proposal

The University of Tennessee  
Arizona State University  
Maryville College

## I. Final Hardware

Our cluster is made up of 5 nodes. The most important of these nodes is a Cray CS-Storm 500NX, containing 8 Nvidia Tesla V100 SXM2 GPUs each drawing 300 watts, and two Intel Xeon Gold 6140's which have a TDP of 140 watts each. The 4 remaining cpu nodes share the same dual Xeons from the GPU node. Every node in the cluster wields 2 Connectx-5 dual port adapter cards, allowing for each node to maintain a direct connection to all other nodes in a mesh network. Each node also possesses a 1TB Intel NVME drive that can be used as scratch space, or in the event of an IO500 run, a parallel filesystem. Our ethernet network switch is a Netgear 16 port unmanaged ethernet switch.

The only hardware that changed from the original proposal was our login node. While the login node is still an Intel NUC, the model was changed due to the lack of cores on the Crimson Canyon NUC. We are now using last years NUC model (NUC7i5BNH) in order to get 4 cores on our login node, although these cores now lack the AVX512 instruction. This means we will have to cross compile. Our NUC is outfitted with 32 GB of RAM, 1 TB of NVME storage and 5 TB of spinning disk.

We hope our configuration will allow us to draw almost all of our power budget from the lone gpu node while running benchmarks, while giving us enough leftover power to idle/suspend the remaining hardware.

## II. Final Software

Centos 7.6 runs clusterwide as our operating system. It was decided upon mostly for driver compatibility. We plan on running the latest version of the NVIDIA drivers, and Mellanox OFED. While the login node is diskful, the 5 compute nodes are managed by xCAT to be diskless. Running diskless compute nodes gives us a more reproducible environment. In the Slurm will be run on our cluster as the scheduler, with our login node serving as the controller.

We ran into a bug with mpirun in the 19.11 version of Slurm, so we had to downgrade to the 19.05 version. We have access to the following compilers on our cluster: PGI, Intel, and GNU. We believe the Intel compiler will be the compiler of choice for cpu only applications, while the PGI compiler will be used on gpu applications.