

DERECHO - NWSC UPGRADE

ALLOCATION PLANS, HARDWARE,
SOFTWARE, AND STORAGE
SPECIFICATIONS

DERECHO: ALLOCATION PLANS – JUNE 2022

- CPU ALLOCATIONS

- Allocations and processes similar to Cheyenne
- Requests in core-hours
- *Estimate Derecho need by multiplying Cheyenne cost x 0.77*

- GPU ALLOCATIONS

- Derecho GPU nodes treated as separate system for allocations purposes
- General content of requests largely the same
- Requests made in “GPU-hours”
- *# of GPUs * wallclock time*
- Base estimates on reasonable performance extrapolation from other GPU system - such as Casper

- UW Share – 13%

- *318 million core-hours / year - ~2x Cheyenne core-hours*
- *328,000 GPU-hours / year*

DERECHO: HARDWARE ENVIRONMENT

- 19.87-petaflops HPE Cray XE System
- 2,570 compute nodes total: 2,488 Homogeneous Compute and 82 Heterogeneous (GPU) nodes
 - Homogeneous nodes have 2x AMD 64-core 2.6 GHz Milan (7H12) CPUs
 - Heterogeneous (GPU) nodes have 1x AMD 64-core 2.6 GHz Milan (7H12) CPU and 4x NVIDIA 1.41 GHz SMX4 A100 Redstone GPUs with 40GiB HBM2 memory and a 600 GB/s NVLink GPU interconnect
- Homogeneous compute nodes have 256GB and the Heterogeneous (GPU) nodes have 512GB of DDR4-3200 memory
- 692 terabytes (TB) of total memory
- HPE Cray Slingshot™ (v11) high-speed interconnect in a Dragonfly topology
 - Homogeneous compute nodes have one Slingshot injection port and the GPU nodes have 4 Slingshot injection ports per node
 - Slingshot bandwidth is 200 Gb/sec per port per direction
 - Slingshot MPI latency is 1.7-2.6 usec
- 8 login nodes, each with 512 GB DDR4-3200 memory
 - six nodes with 2x AMD 2.2 GHz 64-core Rome CPUs
 - two nodes with 2x AMD 2.2 GHz 64-core Rome CPUs and 2x NVIDIA V100 GPUs

DERECHO: SOFTWARE ENVIRONMENT

- Operating System is Cray Linux Environment (CLE), a tuned version of SUSE Linux
- Altair Accelerator Plus scheduler with PBS Professional Workload Manager
- Support for Docker containers, Singularity containers, and containers that support the Open Container Initiative standard.
- Cray Programming Environment (CPE), support for OpenMP 4.5 and 5.0, and MPI v3.1
- CrayPAT: Cray Performance Analysis Tool
- NVIDIA (formerly PGI) Compiler Environment
- Intel Parallel Studio XE compiler suite
- Cray Lustre File System (based on 2.12 LTS)

DERECHO: STORAGE & FILE SYSTEM

- Six HPE Cray ClusterStor E1000 systems
- 60 petabytes of usable file system space (can be expanded to 120 petabytes by exercising options)
- 300 GB per second aggregate I/O bandwidth to/from the NWSC-3 system
- 5,088 × 16-TB drives
- 40TB SSD for Lustre file system metadata
- Two metadata management units (MDU) exporting four MDTs (one MDT exported per one MDS), configured in highly available storage pairs.
- Cray Lustre Parallel File System