

Derecho supercomputer

Derecho, the new supercomputer NCAR installed in 2023, features 2,488 compute nodes with 128 AMD Milan cores per node and 82 nodes with four NVIDIA A100 GPUs each. The HPE Cray EX cluster is a 19.87-petaflops system that is expected to deliver about 3.5 times the scientific throughput of the Cheyenne system. Additional hardware details are available below.



See the following pages for user documentation that is relevant to all NCAR systems (compiling code, environment module basics, managing allocations), and the menu on the right of your screen for system-specific information.

- [Getting started with NCAR systems](#)
- [New user orientation](#)
- [User support](#)

Additional Derecho documentation is in development.

Estimating Derecho allocation needs

Derecho users can expect to see a **1.3x** improvement over the Cheyenne system's performance on a core-for-core basis. Therefore, to estimate how many CPU core-hours will be needed for a project on Derecho, multiply the total for a Cheyenne project by **0.77**.

When requesting an allocation for Derecho GPU nodes, please make your request in terms of GPU-hours (number of GPUs used x wallclock hours). We encourage researchers to estimate GPU-hour needs by making test/benchmark runs on Casper GPUs, but will accept estimates based on runs on comparable non-NCAR, GPU-based systems.

Derecho hardware

323,712 processor cores	3rd Gen AMD EPYC™ 7763 Milan processors
2,488 CPU-only computation nodes	Dual-socket nodes, 64 cores per socket 256 GB DDR4 memory per node
82 GPU nodes	Single-socket nodes, 64 cores per socket 512 GB DDR4 memory per node 4 NVIDIA 1.41 GHz A100 Tensor Core GPUs per node 600 GB/s NVIDIA NVLink GPU interconnect
328 total A100 GPUs	40GB HBM2 memory per GPU 600 GB/s NVIDIA NVLink GPU interconnect
6 CPU login nodes	Dual-socket nodes with AMD EPYC™ 7763 Milan CPUs 64 cores per socket 512 GB DDR4-3200 memory
2 GPU development and testing nodes	Dual-socket nodes with AMD EPYC™ 7543 Milan CPUs 32 cores per socket 2 NVIDIA 1.41 GHz A100 Tensor Core GPUs per node 512 GB DDR4-3200 memory
692 TB total system memory	637 TB DDR4 memory on 2,488 CPU nodes 42 TB DDR4 memory on 82 GPU nodes 13 TB HBM2 memory on 82 GPU nodes

HPE Slingshot v11 high-speed interconnect	Dragonfly topology, 200 Gb/sec per port per direction 1.7-2.6 usec MPI latency CPU-only nodes - one Slingshot injection port GPU nodes - 4 Slingshot injection ports per node
~3.5 times Cheyenne computational capacity	Comparison based on the relative performance of CISL's High Performance Computing Benchmarks run on each system.
> 3.5 times Cheyenne peak performance	19.87 peak petaflops (vs 5.34)