Cheyenne supercomputer

Cheyenne is a 5.34-petaflops, high-performance computer built for NCAR by SGI. The system was released for production work on January 12, 2017.

An SGI ICE XA Cluster, the Cheyenne supercomputer features 145,152 Intel Xeon processor cores in 4,032 dual-socket nodes (36 cores/node) and 313 TB of total memory.

Cheyenne's login nodes give users access to the GLADE shared-disk resource and other storage systems.

Data storage components provided by DataDirect Networks (DDN) give the GLADE system a total usable capacity of 38 PB. The DDN system transfers data at the rate of 200 GBps, more than twice as fast as the previous file system's rate of 90 GBps.

Go to "Quick start on Cheyenne"

To log in, start your terminal or Secure Shell client and run an **ssh** command as shown here:

```
ssh -X username@system_name.ucar.edu
OR
ssh -X username@system_name.hpc.ucar.edu
```

Some users (particularly on Macs) need to use **-Y** instead of **-X** when calling SSH to enable X11 forwarding.

You can use this shorter command if your username for the system is the same as your username on your local computer:

```
ssh -X system_name.ucar.edu
OR
ssh -X system_system_name.hpc.ucar.edu
```

After running the ssh command, you will be asked to authenticate to finish logging in.

Page contents

Hardware

145,152 processor cores	2.3-GHz Intel Xeon E5-2697V4 (Broadwell) processors 16 flops per clock
4,032 computation nodes	Dual-socket nodes, 18 cores per socket
6 login nodes	Dual-socket nodes, 18 cores per socket 256 GB memory/node
313 TB total system memory	64 GB/node on 3,168 nodes, DDR4-2400 128 GB/node on 864 nodes, DDR4-2400
Mellanox EDR InfiniBand high-speed interconnect	Partial 9D Enhanced Hypercube single-plane interconnect topology Bandwidth: 25 GBps bidirectional per link Latency: MPI ping-pong < 1 µs; hardware link 130 ns
3 times Yellowstone computational capacity	Comparison based on the relative performance of CISL High Performance Computing Benchmarks run on each system.
> 3.5 times Yellowstone peak performance	5.34 peak petaflops (vs. 1.504)

Estimating core-hours needed

Cheyenne allocations are made in core-hours. The recommended method for estimating your resource needs for an allocation request is to perform benchmark runs. Some guidance is provided here.

The core-hours used for a job are calculated by multiplying the number of processor cores used by the wall-clock duration in hours. Cheyenne core-hour calculations should assume that all jobs will run in the regular queue and that they are charged for use of all 36 cores on each node.