

## High-Performance Computing (HPC) Resources at CSU

CSU is extremely well positioned with HPC infrastructure to support research and related activities, as described below. Additional detail is available at <http://hpc.colostate.edu>, where this document can be found.

### A. Central HPC Cray XE6

Central IT at CSU operates, maintains and supports a 2,560 core Cray XE6, with 32 GB of RAM per node, utilizing the high-speed Gemini interconnect. Scratch storage consists of 32 TB of fast, parallel lustre storage, augmented by approximately 100 TB of expandable user storage space. The system is connected at 10 Gbps to the ultrahigh-speed research DMZ LAN. The ~25 TFLOP system has been highly stable and extremely productive. There is ample available capacity on the system for relatively modest jobs. More information is available on the ISTE C Cray website [istec.colostate.edu/activities/hpc/cray](http://istec.colostate.edu/activities/hpc/cray).



### B. Joint NSF-funded Summit System

Additionally, CSU and the University of Colorado at Boulder have been awarded an NSF grant in the total amount of \$3.5 million. Summit, the HPC system purchased under this award, went into operation on February 2, 2017. Summit is shared between the two universities, and jobs up to the full size of the system are supported. The specifications for the original system are given in Table 1 below.

Table 1 Summit HPC System Specifications						
Node type	Count (nodes)	Proc. Freq. (GHz)	# cores per node / total	Memory per core / node (GB)	TFLOPS per node / aggregate (DP)	Co-Processor
Xeon Haswell	380	2.3	24 / 9,120	5.3 / 128	0.9 / 338	
Xeon KnL-F	20	1.5	72 / 1,440	0.2/16 (MCDRAM) 5.3 / 384 (DDR4)	3 / 60	
GPU	10	0.6	9,984 / 99,840	0.005 / 48 (GDDR5)	6 / 60	2x Nvidia K80 cards per node
HiMem	5	2.1	48 / 240	43 / 2,048	2 / 10	
Total	415				468	

The system uses the new OmniPath interconnect, and includes 1 Petabyte of scratch storage. This is a very large, ultra-modern heterogeneous HPC system, with ample capacity for extremely large simulations. Accounts and allocations are available through the department of Academic Computing and Networking Services (ACNS) at <http://hpc.colostate.edu>.

### E. Condominium Node Buy-in

Since Summit went live, we have held two rounds of solicitation, where users have elected to buy into the Summit system under the “condominium node” model. Here a user purchases a compatible node of type specified by the system administrator to be added to the Summit system. The user purchases just the basic node, and the shared hardware (interconnection fabric, rack, power distribution unit, installation, etc.) is subsidized by CSU ACNS. The model provides a “take back” guarantee; if at any time, the user is dissatisfied with the service, the node may be repossessed by the user. The user is given a full allocation of time for the node(s) purchased, matched by an equal allocation of time on the general Summit system.

Two rounds of condo node buy-ins have been held, with the result that the system capacity has been expanded on the order of twenty-five percent. No additional condo node buy in solicitations for adding to Summit are anticipated, as the hardware is getting too old. However, there is still ample capacity on Summit for CSU users, who can use the system for free.

CU and CSU are contemplating jointly writing another MRI grant proposal to the NSF for a next generation system, into which researchers may buy under the condo node model. Please stay tuned for this eventuality.

### **C. HPC Courses**

CSU teaches five courses in HPC, several specifically at the graduate level targeted toward general users. The courses are oriented to teach students programming using shared memory parallelism (OpenMP) and distributed memory parallelism (MPI, OpenMPI). Large (even huge) real-world data sets are used emphasizing High-Throughput Computing.

### **D. Networking Infrastructure**

Local Area Network (LAN) – CSU currently operates a fully redundant, self-healing 10 Gbps core backbone network, with a 40 Gbps ultrahigh-speed Research DMZ LAN. Primary links to major buildings are 10 Gbps, and all major buildings are well covered with 802.11ac Wi-Fi. Devices and systems needing ultrahigh-capacity are configured on the Research DMZ, providing the bandwidth they need, as well as taking a high volume of traffic off of the building network, rendering it usable by users.

Wide-Area Network - CSU obtains excellent Internet access from the FRGP in Denver, Colorado, a collaboration that began in 2001. Access includes redundant links to two commodity providers: TeliaSonera and Level3; and various research networks: Internet2, WRN, and ESnet. CSU is currently connected to the FRGP at 30 Gbps, with a 10 Gbps wave dedicated to research. The FRGP is near completion of an upgrade to 100 Gig transport utilizing Adva wave-division multiplexing optical equipment, providing additional headroom.

More than ample capacity exists in both the LAN and the WAN to support the most stringent needs.

### **F. File Storage and Sharing**

CSU has installed an expandable Data Transfer Node on its Research DMZ, and provides Globus for granular file sharing among individual faculty members, in addition to very high speed, parallel file transfer. Faculty may buy into the system, if they wish. The system is operated by central IT, with the requisite security and privacy mechanisms in place to exceed typical federal grant requirements. Those desiring more information should contact ACNS.

### **G. Shared Digital Repository**

Due to its extensive experience operating a digital repository, CSU hosts the Shared Services DSpace Digital Repository, encompassing seven campuses in Colorado. That activity is organized specifically to share expertise, realize economies of scale, and take advantage of geographically distributed data centers for storage and preservation. It is this repository that will be used to ingest, curate, expose and preserve scholarly communications and associated data sets to comply with data management requirements.

### **CyberTeams Support**

Under the leadership of the University of Colorado and with participation by the University of Utah, CSU participates in the NSF CyberTeams grant. This grant provides “cradle to grave” support for research groups at CSU, CU and UU, for research data flows, such as:

- Data security and privacy.

- Pre-processing, e.g. input files preparation.
- HPC support for large-scale analysis.
- Post-processing, e.g. scientific visualization, and
- Data Curation and Preservation.

Generally, the University of Utah provides primary consultative capacity on IT Security and Privacy, CU supplies consultation on data transfers, data privacy, and data security. Also, the CSU Libraries provides consultation on preservation.

### **Glossary of Acronyms**

ACNS – Academic Computing and Networking Services (central support for HPC at CSU)

BiSON – Bi-State Optical Network

CI – Cyberinfrastructure

CSU– Colorado State University in Fort Collins, Colorado

DMZ – Demilitarized Zone, an ultrahigh-speed portion of the LAN dedicated to research

DTN – Data Transfer Node

FRGP – Front Range GigaPoP

GPU – Graphics Processing Unit

HPC – High Performance Computing

HTC – High Throughput Computing

I2 – Internet2

Gbps – Gigabits Per Second

LAN – Local Area Network

TFLOP – A metric of computing performance, 1 Trillion ( $10^{12}$ ) floating-point operations per second

WAN – Wide Area Network