

High-Performance Computing (HPC) Resources at CSU

CSU is extremely well positioned with HPC infrastructure to support research and related activities, as described below.

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.



B. Joint NSF-funded System

Additionally, CSU and the University of Colorado at Boulder have been awarded an NSF grant in the total amount of \$4 million, and are in the process of accepting bids due Nov.16, 2015 for a new very high performance HPC system that is to be installed by April 2016. The system will be shared between the two universities, and jobs up to the full size of the system will be supported. The system will meet or exceed the following 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	376	2.3	24 / 9,024	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	411				468	

The intent is to use the new OmniPath interconnect, and 1 Petabyte of storage will exist. This will be a very large, ultra-modern HPC system, with ample capacity for extremely large simulations.

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). 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.

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.

E. File Storage and Sharing

CSU has installed a large and expandable Research Data Storage System that is located behind the campus firewall to comply with federal and university policies. This storage is provided through a Dell Isilon Gen 6 configuration consisting of four H500 nodes with 324 TB Usable space. This system is incrementally expandable by adding two nodes for an additional 216 TB of storage. The system leverages Isilon's multi-protocol support allowing files to be dragged and dropped from users' computers, has snapshot capabilities. This system is operated professionally by central IT staff in ACNS. However, purchase and management of User quotas are managed flexibly in the distributed environment through the smart quota software, where storage may be provisioned to users at their college or departmental level, or by themselves. The system is operated in cost recovery mode at a rate of \$70/TB-yr. Users interested in this service should contact Mr. Dave Hoffman, dave.hoffman@colostate.edu.

CSU also has installed a centrally deployed, managed and operated file storage system for Controlled Unclassified Information, as required in some federal contracts. This is a new service that is currently being deployed, and for which a financial cost recovery model is under development. Users interested in this service should contact Mr. Dave Hoffman, dave.hoffman@colostate.edu.

E. Shared Digital Repository

Due to its extensive experience operating a digital repository, CSU hosts the Shared Services 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.

Glossary of Acronyms

ACNS – Academic Computing and Networking Services

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 (10¹²) floating-point operations per second

WAN – Wide Area Network

Version updated January 31, 2018 by Dave Hoffman and Patrick J. Burns