Facilities, Equipment and Other Resources

Shared supercomputer between CU-Boulder and CSU - regional access

The Summit supercomputer, funded by NSF under Grant No. AC- 1532236, is currently in acceptance testing. The system has peak performance of over 400 TFLOPS. The 380 general compute nodes have two Intel Haswell CPUs with 12 cores each, 128 GB of RAM and a local SSD. Additionally, the system has 10 GPU nodes containing two NVIDIA K80 GPUs, 5 high-memory nodes with about 1 TB of main memory and in a second deployment in December 2016 20 Xeon Phi ("Knight's Landing") nodes with 72 real cores supporting 288 threads for development and benchmarking. All nodes are connected through a high-performance network based on Intel Omni-Path with a bandwidth of 100 Gb/s and a latency of 0.4 microseconds. 1 PB of high-performance storage is provided using the IBM GPFS file system. This system is available to CU-Boulder researchers and collaborators, as well as 10% of cycles are provided to members of the Rocky Mountain Advanced Computing Consortium.

Expertise

Each of the institutions partnered for this proposal have is a small groups of computational scientists, high-performance computing specialists, and system administrators with the mission to provide leadership in developing, deploying, and operating an integrated cyberinfrastructure consisting of high-performance computing, storage, and high speed networking that supports and encourages research, collaboration and discovery. The groups also contributes to the educational mission of the university by providing workshops and training on cyberinfrastructure related topics.

Facilities at CSU

Local and Wide Area Networks

We have recently deployed a "Research LAN" interconnected at 40 Gbps. We have a dedicated 20 Gbps connection for internet access, and an additional, and a separate 10 Gbps wave dedicated to research, to the FRGP our regional optical network provider for internet access, to which we have a fiber connection. This provides more than adequate capacity for our applications. Note that the core infrastructure at the FRGP is capable of 100 Gbps, and we can easily upgrade the current research wave to 100 Gbps, if needed.
Networking tools - Note on the diagram that we have perfSONAR deployed in several locations, and have deployed NetSec for the Research LAN.

Shared digital repository - CSU operates, supports, and maintains a shared Digital Repository for nine libraries in Colorado. It is this that we use for ingesting, making discoverable and accessible, and preserving scholarly communications and research data sets. Indeed, our data management plan relies on placing relevant materials from this activity there (see our Data Management Plan). Note that our current DigiTool product from Ex Libris is planned to be upgraded to either the Dspace or Fedora/Islandora open source product by the end of Q2 2015.

Data Transfer Node

Note that we have implemented with internal funding a Data Transfer Node, including Globus, that supports large file storage and transfer. This system is now directly connected to the Research LAN, and also supports big data analysis on our Cray XT6m system.

Shared digital repository

CSU operates, supports, and maintains a shared Digital Repository for nine libraries in Colorado based on DSpace. It is this that we use for ingesting, making discoverable and accessible, and preserving scholarly communications and research data sets. Indeed, our data management plan relies on placing relevant materials from this activity there (see our Data Management Plan).

University of Colorado Boulder

Networking

The current CU Boulder network is a 40 Gbps fiber core with Cat 5 or higher wiring throughout campus. RC has created an 80 Gbps Science-DMZ to connect the Janus supercomputer to storage and to bring individual dedicated 10 Gbps circuits to various locations as needed. CU Boulder participates in I2 (the Internet 2 higher education, government, and vendor research computing con- sortium) and is an active member of the Front-Range gigapop and other networks. RC has started to provide campus researchers with a leading edge network that meets their needs and facilitates collaboration, high performance data exchange, access to co-location facilities, remote mounts to storage, and real-time communications.

File Transfer

For moving large amounts of data Research Computing has several nodes dedicated to GridFTP file transfer. RC's GridFTP servers support both the Globus Connect web environment and basic GridFTP via the command line.
OIT also offers a file transfer service with a web interface, which provides a good way to transfer files to collaborators. Files are uploaded to a server and a link to download the file can be emailed to an on or off-campus user.

Compute

- RC operates the joint CSU-CU-Boulder Summit supercomputer, funded by NSF under Grant No. AC- 1532236
- RC provides a small "HT cluster" for jobs requiring more RAM or longer runtimes than are available on the Janus supercomputer nodes. The cluster consists of a Dell M1000e chassis filled with 16 blades. Each blade (i.e., node) has 96 GB RAM, 12 physical cores (24 w/hyperthreading) and 2 TB of local disk. This system is intended for single-node jobs as the nodes do not have high-speed, low-latency network interconnects. The CPUs are Intel Xeon X5660, 6 cores, 2.8Ghz, 12M Cache.
- RC provides three "Analytics and Visualization" nodes. One node has 512 GB RAM, 2 TB local disk, 32 physical cores (64 w/hyperthreading) and four Intel Xeon X7550 processors. Two nodes provide 1 TB RAM, 16 TB local disk, 40 physical cores (80 w/hyperthreading), and four Intel Xeon E7-L8867 processors.

Storage

Each researcher using the computational resources at CU Boulder has a home directory with 2GB and a project space consisting of 250 GB of storage. Additional storage is provided as part of a storage condominium at a cost of $100 per TB for single copy storage. Tape and HSM are additional storage options that are available for archive data.

PetaLibrary

The two main categories of service offered to customers of the PetaLibrary are Active storage for data that needs to be accessed frequently and Archive storage for data that is accessed infrequently. Active data is always stored on disk and is accessible to researchers on compute resources managed by RC. Archive storage consists of a two level hierarchical storage management (HSM) solution, with disk storage for data that is more likely to be accessed and tape for data that is less likely to be accessed frequently. The cost for the research is $100/TB/year for disk and $20/TB/year for tape.

University of Utah

Data Center

The University of Utah opened a state of the art, LEED Gold certified Data Center with 74,000 square feet of space in the spring of 2012, and has developed over 30,000 square feet in the initial phase for both enterprise computing (administrative, academic and medical uses) and high performance computing and research. The high performance computing and research component is located in a dedicated room with 1.15 MW of power and cooling available. Of this total power, 150 kW is Tier 3 Power with redundant UPS and diesel generator backup for at least three days. The remaining 1 MW, which supports the computational clusters, has power

conditioning and battery backup. The University manages the building with 24/7 support staff, providing both monitoring and access/security presence.

Networking

The University of Utah maintains a fully redundant campus backbone network that serves the administrative, academic, and research needs of the main campus and the needs of the University of Utah Hospital and Clinic infrastructure. The University maintains multiple security zones in order to balance the data security and compliance requirements of the various constituencies. Multiple layers of security controls apply to various zones. The campus backbone supports technologies such as Multi-Protocol Label Switching (MPLS), Virtual LANs (VLANs), Virtual Routing instances, and both air-blown fiber, and conventional optical fiber. The campus backbone fully supports both IPv4 and IPv6 routing (BGP, EIGRP, OSPF, and OSPFv3) and multicast routing. The campus backbone supports multiple 10-Gbps and 40-Gbps backbone links today, and the campus uplinks to the Utah Education Network (UEN) via 100Gbps. The University of Utah campus connects to the new off-campus Downtown Data Center and to its upstream Internet providers by the University/UEN metropolitan optical (DWDM) ring that extends over the greater Salt Lake City area. This Metro Optical network supports multiple 10-Gbps and 100-Gbps wavelengths, up to a capacity of 8.8 Terabits per second. UEN currently connects via this optical network to the Internet2 Network at its intermountain regional node in western Salt Lake City with a 100-Gbps link and supports the University of Utah with multiple 10-Gbps connections and a 100-Gbps connection.

Computing

The Center for High Performance Computing supports much of the computational and Big Data needs on the campus. Most large domain science computations occur on the CHPC infrastructure. CHPC has a mixture of general CHPC nodes, with access to these resources made available at no cost to the users, and owner nodes. The time on the general CHPC resources is awarded by an internal allocation process. Owner nodes are available, with a pre-emption policy, for general usage when not being used by the owner group. CHPC currently manages six compute clusters:

- APT Emulab cluster, provisioned 2014. Currently provides dynamic access to 64 dual socket Intel Sandybridge nodes (1024 total cores).
- Kingspeak, started in August 2013 and still growing. Currently 292 compute nodes, in a mixture of Intel Sandybridge, Ivybridge and Haswell architecture, for a total of 5680 cores. Has Mellanox FDR Infiniband interconnect.
- Ember (2010). 162 Intel Westmere nodes (2304 cores) along with 11 GPU nodes. Has Mellanox QDR Infiniband interconnect

- Ash, built in 2010, expanded in both 2014 and 2015. A total of 253 Intel Westmere nodes, 164 Intel Ivybridge nodes, and 48 Intel Haswell nodes (7468 cores). Has Infiniband interconnect.
- Apexarch, a small cluster with 16 nodes (152 cores) that is used for HIPAA and other protected applications.
- Lonepeak, a mixed cluster, with some large memory nodes, for a total of 141 nodes (1376 cores).

In addition, CHPC supports two VM farms, one capable of handling restricted data, which host a variety of applications, including web services, for our users. The Center has a staff of computer professionals with a great deal of expertise in the operation of large computer systems and a scientific user support team available for software installation, consultation on usage of the resources, and collaboration on projects.

Storage

CHPC also supports various scratch, group and home directory file systems with a total of over 14 PB of usable space supported. The infrastructure interconnects at bandwidths ranging from 10-Gbps and 40-Gbps for Ethernet to 56 Gbps for Infinband. This infrastructure connects to the campus backbone at 10-Gbps and the Science DMZ at 100-Gbps with over a dozen data transfer nodes deployed at either 10-Gbps or 40-Gbps.

Protected Data

CHPC also operates a protected environment (PE), as briefly mentioned, consisting of both a linux cluster, windows server, and a virtual machine (VM) farm along with associated storage for researchers with sponsored research projects and work with data that is sensitive in nature. The resources and processes have been reviewed and vetted by the Information Security Office and the Compliance Office as being an appropriate place to work with Protected Health Information (PHI). Servers and disks in the PE are stored in locked cabinets that are accessible to even fewer staff. The rack containing the hardware is locked, and the devices containing sensitive data also have locked bezels, preventing local console access. CHPC has isolated this PE and uses a network protected logical partition that provides research groups specific access to individual data sets.