University at Buffalo Center for Computational Research

The following is a short and long description of CCR Facilities for use in proposals, reports, and presentations. If desired, a letter of support for your proposal can also be written. In addition to the computing and visualization resources listed below, CCR staff members, who have extensive expertise in high-performance computing, high-capacity data storage, advanced database design, custom software development, grid computing, bioinformatics support, and scientific visualization, are available to facilitate your research.

At a minimum, faculty using CCR resources are required to:

1. **Acknowledge CCR support in publications**,
2. Include **a percent credit to CCR (at least 5%)** on the Sponsored Programs Approval Form (though this can be larger for grants requiring significant CCR resources). NOTE: Clicking the "use CCR" box on the Sponsored Programs Approval Form is not the same as including a percent credit to CCR.

Recommended CCR Acknowledgment: *This work was performed in part at the University at Buffalo's Center for Computational Research.*

Want your jobs to go faster?

While there is no cost for UB faculty groups to use CCR compute resources, queue priority is based on contribution level. In order to receive a priority boost in your jobs, you must have a **funded grant** that meets the criteria above and **includes direct funds to CCR.** At minimum, you should budget $2,600 per year (40 hours/year of CCR support staff time at the SPA approved rate of $65/hour for typical projects - higher amounts as appropriate for projects that are expected to more heavily utilize CCR services). Boosting your group's priority will substantially increase your job throughput. Please notify CCR Help with your grant information so we can provide you with the boost in queue priority.

SHORT DESCRIPTION:

The Center for Computational Research (http://buffalo.edu/ccr), a leading academic supercomputing facility, maintains a high-performance computing environment, high-end visualization laboratories, and support staff with expertise in scientific computing, parallel processing, and data analytics. The Center's extensive computing facilities, which are housed in a state-of-the-art 4000 sq ft machine room, include a generally accessible (to all UB researchers) Linux cluster with more than 8000 processor cores and QDR Infiniband, a subset (32) of which contain (64) NVidia Tesla M2050 "Fermi" graphics processing units (GPUs). Industrial partners of the University have access to an additional cluster with more than 3400 processor cores and FDR Infiniband. The Center maintains a 3PB IBM GPFS high-performance parallel file system plus a 1PB EMC^2 Isilon shared network attached filesystem. The computer visualization laboratory features a tiled display wall, and a VisDuo passive stereo system. A leading academic supercomputing facility, CCR has more than 250 TFlop/s of peak performance compute capacity. CCR additionally hosts a number of clusters and specialized storage devices for various specific

departments, projects, and collaborations, researchers interested in hosting services should contact CCR staff.

In addition to its computing and visualization resources, CCR has a support staff consisting of computational scientists, programmers, and database administrators with expertise in all areas of computing, including scientific and parallel computing, (big) data analytics, custom software development, database engineering, and scientific visualization.

MORE DETAILED DESCRIPTION:
CCR's resources include the following.

"Rush" x86_64 Linux Cluster: CCR's x86_64 Linux cluster consists of 372 Dell C6100 servers, each with two Intel "Westmere" Xeon 2.40GHz (E5645) processors and 48GB of memory, 128 IBM iDataPlex dx360 M2 servers, each of which has two Intel "Nehalem" Xeon 2.26GHz (L5520) processors and 24GB of memory, 128 Dell C6100 servers containing two Intel "Westmere" Xeon 2.13GHz (L5630) processors and 24GB of memory, and 34 Dell R620 servers each with two Intel "Sandy Bridge" 2.26GHz Xeon (E5-2660) processors. All are interconnected with QDR Infiniband and gigabit Ethernet. A sub-cluster of 32 servers (Dell C6100 servers with two Intel "Westmere" Xeon 2.66GHz (X5650) six-core processors, 48GB of memory) is outfitted with two NVidia Tesla M2050 "Fermi" graphics processing units (GPUs), and solid-state drives for data-intensive operations. A single server (Dell R910) with four Intel "Beckton" Xeon 2.0GHz (X7550) eight-core processors has 256GB of main memory and a local pool of 1.2TB of solid-state hard drives for extreme needs of memory and input/output. Additional "fat" memory nodes are available with 32-cores each and either 256GB (8 with Intel "Westmere" Xeon (E7-4830) processors, 8 with AMD "Magny-Cours" Opteron (6132HE) 2.2GHz processors) or 512GB of memory (2 with Intel "Westmere" Xeon (E7-4830) processors) and 4TB of local disk. In addition to the high-performance low-latency Infiniband network, these servers are interconnected with a 10 gigabit Ethernet (10gigE) core network from Arista (6508 switch), with each server connected to gigabit Ethernet leaf switches, uplinked by dual 10 gigabit Ethernet links to the core. The peak performance of this cluster overall is approximately 100 TFlop/s.

Industrial Users Cluster: This cluster consists of 216 HP SL230 Gen8 servers of which 144 are "Parallel" compute nodes with an FDR InfiniBand interconnect. The remaining 72 nodes are "Serial" nodes with GigE connections. The nodes consist of 2 Intel "Ivy Bridge" Xeon 2.6GHz (E5-2650V2) 8-core processors (total of 3,456 cores), 64GB of memory and 500GB of local scratch. This cluster is primarily for industrial partners and the idle CPU cycles are made available to CCR users with the understanding that their jobs will be preempted if an industrial user job is submitted. The peak performance of this cluster overall is approximately 70 TFlop/s.

Faculty/Departmental Clusters: CCR hosts a variety of clusters for researchers in various departments. Such faculty/departmental clusters are maintained similarly to the general production systems with a common operating environment. Access to these resources is generally reserved to the faculty/department that procured them, but also includes a pre-emptible scavenger partition/queue that allows the general research community to leverage them when underutilized by the faculty/department. Currently these clusters number more than 325 nodes of various

configurations (some including GPUs) across more than 16 faculty/departments. All such clusters are also integrated within the core CCR 10gigE Ethernet network and storage systems. Frequently such systems are used for specialized instruction in HPC topics, or for more general disciplinary use (e.g., a small cluster is routinely used for chemistry instruction on a shared enterprise WebMO service).

IBM GPFS Parallel Filesystem: The GPFS storage solution consists of Lenovo servers and provides 3PB of available storage. This storage is CCR's high performance parallel file system. With 40GigE connections to the core Arista network, it provides I/O performance in excess of 30 GigaBytes per second sustained in tests. Designed for high performance and concurrent access, CCR's GPFS is primarily intended for generation and analyses of large quantities of short-lived data (scratch usage).

EMC^2 Isilon Filesystem: The EMC^2 Isilon storage system consists of 1PB of usable storage in a hierarchical storage pool of three X410 storage servers and six NL410 storage servers, connected to the CCR core network with two 10GigE links per server. This storage system serves as the high reliability core storage for home and projects usage. The storage is designed to tolerate simultaneous failures, helping assure the 24x7x365 availability of the Center's primary storage.

Tiled-Display Wall: This 12'x7' device was assembled for the purpose of allowing scientific investigations of high-resolution images by teams of scientists working in a comfortable setting. In addition, the tiled-display wall is ideal for many of CCR's urban planning projects. The tiled-display wall is back-projected by 9 1920x1080 projectors arranged in a 3x3 matrix providing 15.2 megapixels of resolution. A single custom-built server containing dual quad-core Intel Xeon processors, 48 GB of memory and 5 1024MB NVidia Quadro 4000 Tesla-based graphics cards, drives the 9 projectors. Additionally, the individual tiles are edge blended in order to provide a large, seamless, high-resolution image. The machine is housed in the base of the tiled display wall, providing a self-contained display device.

Passive Stereo Display: The VisDuo is a ceiling mounted, dual projector, and passive stereo display, used for viewing complex 3D environments, molecular structures, and medical simulations. The stereo effect is generated by each projector producing images for one eye, whose output is polarized by special filters and the resulting image is viewed on a custom polarization preserving screen. Users can view the resulting 3D imagery by wearing lightweight polarizing glasses. This system is powered by a custom workstation with a 3.4 GHz Intel Xeon Processor, 2 GB of RAM and an NVidia Geforce 7800 GTX graphics card.

Staff: CCR maintains a support staff of 13 people, including 3 Ph.D. level computational scientists, an HPC support specialist 5 systems analysts/administrators, a database administrator/software engineer, 2 visualization specialists/software engineers, and 1 administrative support staff.

External Network: In addition to high-end computing and visualization resources, high-speed communication is also available. The University at Buffalo is an Internet2 member.