

Cluster Partnership Program

This document defines the CRC's Cluster Partnership Program (CPP) for ND Faculty regarding system procurement, administration, and support. Under this program the CRC provides specification/acquisition expertise, space (in an appropriate and secure operating environment), utilities infrastructure (e.g., power, cooling), system administration, and server/hardware support (warranty work coordination).

Given the CRC's provision of infrastructure and services, our partner researchers are able to leverage their capital resources to purchase approximately 30% more compute nodes while significantly reducing the overhead of internal system administration. In return for infrastructure and services provided by the CRC, when partner compute resources are not being used by the partner they are made available to other "trusted and experienced" users within the University's computational community.

Partner Program Details

Partner program hardware options are compatible with the CRC's general access hardware resources. Compatible hardware allows CRC operations staff to effectively support a large number of systems—since the technical efforts required to manage the compute node clusters increases only modestly with additional, compatible hardware.

Multiple compatible high-performance computing (HPC) environments are available for partners based on the CRC's general access HPC environments. Examples include:

- Distributed memory Linux cluster environment for high throughput computing with gigabit Ethernet interconnections.
- Distributed memory Linux cluster for highly parallel computing with ten gigabit Ethernet or Infiniband interconnections and high end RAM/Disk/CPU components.
- Shared memory Linux systems—up to 1TB of RAM and multiple CPUs.
- Accelerator based systems with NVIDIA GPUs or Intel Xeon Phi cards

CRC personnel annually test and negotiate pricing for equipment from specific "tier-1" vendors (IBM, Dell, or HP). The partner cost is the actual cost of a (rack-mountable) server including the network card, a three year maintenance warranty, switch gear, server rack, power distribution unit (PDU), and cabling. Partners who are able to bring at least 1 full server are welcome to participate. Switch gear, rack, and PDU costs for less than a full rack of servers will be charged in ½ capacity increments. For example a standard rack holds 36 servers, faculty purchasing only 10 servers could be charged for up to ½ rack, ½ switch, etc....

Management of Partner Compute Resources

All systems use the CRC's software stack: Red Hat Linux, Grid Engine, CRC file systems, and faculty-requested applications installed on the CRC's distributed file systems.

Grid Engine (GE) share-based priority scheduling is used to provide equitable access to compute resources. Partners have their own dedicated GE queues to manage their exclusive resources. All access to compute node cycles is through GE; faculty may purchase a dedicated interactive login node or share the general access nodes which provide access to CRC resources.

In return for infrastructure and services provided by the CRC, when partner compute resources are not being utilized they are made available to the "trusted and experienced" users within the University's computational community. If a partner requires his/her previously unused resource back, the guest user jobs are immediately evicted (~5min until the owner job starts). The CRC assumes responsibility to ensure nodes will be returned to use by the faculty without detrimental remnants from prior guest jobs/simulations (the ~5min is for clean-up scripts).

System components are warrantied for 3 years, after which hardware replacement costs are the responsibility of the faculty member. Typically the useful life of a system has been 4-6 years. As equipment ages, the CRC will work with all partners in advance to retire the older system and seek funds for hardware replacement. The CRC cannot support aging systems indefinitely.

Faculty Members Requiring Special Configurations or Exclusive Resource Access

The CRC will support faculty funded systems outside of the Cluster Partnership Program under the following guidelines.

1. Faculty member funds all necessary hardware to include but not limited to servers, racks, network switches, power distribution strips, and cabling.
2. Faculty member pays the CRC for installation, configuration, and startup services.
 - a. The charge will be in 4 hour increments based on annually approved CRC recharge rates for systems engineering.
3. After system startup the faculty member can utilize software services in line with general CRC resources (CRC distributed storage, CRC software stack, etc...) at no charge.
4. Faculty member pays for system diagnostics and maintenance requirements such as verification of failed hardware components, service call coordination, firmware upgrades, and operating system upgrades/patches.
 - a. The charge will be in 1 hour increments based on annually approved CRC recharge rates for systems engineering.