

# HPC Contract

Apr 2007 version 1.0

## Definitions

The object of this Contract is high-performance computing infrastructure, or HPC infrastructure. The rationale behind the Contract is discussed in the accompanying document “HPC Sustainability Plan”.

**Commodity computing and networking infrastructure** is defined as the collection of hardware, software, personnel, and organizational structure to support what have become common activities in any business or organization, including universities. Examples include reading and sending e-mail, searching and retrieving information from the Internet, reading, composing and printing electronic documents, performing calculations with spreadsheets and specialized software such as Matlab, Maple, and Mathematica, managing information for the operation of the organization with databases and enterprise resource planning systems.

**High-performance computing** is any activity involving running research-oriented software on large and complex computer systems for a long time, days or weeks. The structure of the calculations can vary from loosely parallel to very tightly coupled. The defining characteristic of HPC is that trying to run the same computation on a single laptop or desktop computer would take at least a large fraction of a year, sometimes hundreds of years or even longer.

**HPC infrastructure** is defined as the collection of hardware, software, personnel and organizational structure to support academic high-performance computing research. Because the use of HPC in science and engineering research is changing rapidly, the definition deliberately is kept sufficiently vague to respond to this evolution in the next few years. However, we wish to make clear that neither commodity computing and commodity networking infrastructure nor highly specialized equipment for research in computer and network architecture per se are part of HPC infrastructure as it pertains to this Contract.

## The classes of partners

The partners to the Contract who must agree to build, maintain and grow the HPC infrastructure are the following:

1. Individual members of the faculty representing their research group as principal investigators (PIs).
2. The deans, department chairs, and directors (DDD) of units involved in HPC. Deans, Chairs and Directors may ask advice from the governance structures of their

## HPC Contract

respective units, e.g. departmental strategic planning committee and the unit level computer committee, and the faculty. They represent all faculty members in the respective unit involved in HPC.

3. The Chief Information Officer (CIO). The CIO may involve committees and individuals to reach decisions. The CIO represents the interest of the University of Florida as a whole, and therefore represents the interest of the President, the Provost, the Vice President for Health Affairs, the Vice President for IFAS, and the Vice President for Research, in addition to any and all faculty in the university involved in HPC insofar as they are not already individual partners or are partners through a department or college.

Note the any particular researcher may be represented in the Contract in multiple ways.

### The common goals

The partners recognize the value of building a shared HPC infrastructure. A common HPC structure has the following advantages over multiple isolated efforts:

1. A shared HPC infrastructure can be sufficiently large to provide national visibility and competitiveness for the HPC research at UF.
2. A shared HPC infrastructure can be sufficiently large to provide an attractive work environment for the HPC staff so that extremely talented people can be attracted and retained to work in the HPC Center.
3. A shared HPC infrastructure can provide uniform and streamlined access to more complex and more sophisticated equipment than any single research group can provide for itself.
4. A shared HPC infrastructure can provide efficient access to underused capacity, thus providing better overall utilization and thus justification for advanced state-of-the-art resources.

The common goals by themselves do not provide sufficient incentive to enter into a Contract and uphold it. All of the classes of partners are required to make the shared HPC infrastructure work.

The following three sections specify the benefits and responsibilities associated with each class of partner.

### The researchers

Individual faculty members can be partners, and are strongly encouraged to become partners, in the Contract for shared HPC infrastructure. For the Contract to work there must be at least two researcher-class partners and the university administration.

**Benefits.** Any PI, who is a partner in the Contract for shared HPC infrastructure by committing to assume the responsibilities listed below, will obtain the following benefits.

1. Access to a computed part of the shared resource. The portion of the shared resource is proportional to the amount invested, computed as the number of CPUs, and is equal

to **Portion = Investment \* TotalNumberCPUs / TotalCost = Investment / \$1,733.**

The scheduler of the shared resource will treat the investor and his research group with higher priority to provide access to any number of CPUs less than or equal to the portion defined by the size of the investment with the result of guaranteeing that any job that fits in the investment portion will start in a fixed finite time, seconds or minutes, at most hours<sup>1</sup>. Of course, these researchers have in addition access to portions by virtue of investments made by the units, department or college, they belong to and to the entire resources by virtue of the investments made by the CIO. Access to these portions is provided at different, decreasing priorities. The size of the investment portion will be recalculated every time the shared resource is reconfigured to add new resources or to replace resources. After three years, the boost in priority provided by the investment will slowly decrease to reach zero when the investment is ten years old.

2. Support for the HPC resources by skilled and properly paid professionals, who have a long-term relationship with the HPC infrastructure. This staffing is to be contrasted with short-term temporary workers often put in charge of research group resources.
3. Support for the software to be run by the researcher and his or her associates from skilled and knowledgeable professionals. They will help optimize the performance for the researchers, both for running on the hardware and for scheduling the workload of the researcher.

**Responsibilities.** Any researcher who commits to the following responsibilities becomes a partner in the Contract for shared HPC infrastructure and is thereby entitled to all benefits of the Contract listed above.

1. The researcher will participate in the governance of the shared HPC infrastructure by communicating with representatives or by serving on the relevant committees.
2. The researcher will seek funding for his or her research and either include a budget for equipment acquisition to be managed as part of the shared HPC infrastructure or designate that the center portion of returned overhead of some grants go to the HPC Center.

## The academic administration

The college deans, department chairs and the directors (DDD), or their representative can be partners, and are strongly encouraged to become partners, in the Contract for shared HPC infrastructure when there are some faculty members in their respective units who have interest in or are using HPC in their research. For the Contract to be sustainable, the deans, department chairs and directors of units with research faculty active in HPC must be partners, at least at the minimal commitment level specified below.

**Benefits.** Any unit that is a partner in the Contract for shared HPC infrastructure by committing to the responsibilities listed below ensures the following benefits for its faculty members needing to use HPC in their research.

---

<sup>1</sup> This assumes that the PI does not already have a job running in his or her share. Prioritization of jobs belonging to the same research group inside the PI portion follows standard scheduler practices and capabilities.

## HPC Contract

1. The academic administrator can support all faculty members in the unit interested in HPC in a concrete way without having to take on the burden of management and administration of an HPC computing facility.
2. By investing in the campus shared HPC infrastructure, the faculty in the unit is provided with access to and support from a better and more powerful facility than can be built with unit-level resources.
3. A well-run and visible shared HPC infrastructure is an asset to the administrator, the unit and the university in attracting candidates for new positions and in negotiating startup conditions for new hires.

**Responsibilities.** Any unit that commits to the following responsibilities becomes a partner in the Contract for shared HPC infrastructure and is thereby entitled to all benefits of the Contract listed above.

1. The academic administration (DDD) will advertise to faculty the benefits of joining the UF HPC Center. In particular, they will explain HPC membership as an alternative to building of isolated clusters, hiring isolated system administrator personnel to maintain such clusters, or renovating buildings to provide space, power or cooling for such clusters<sup>2</sup>. In addition they will explain the global benefits and efficacy of investments in the framework of the campus-wide effort to build and maintain the UF HPC Center.
2. The academic administration also should ensure that the faculty members in the unit participate in the governance structure and process of the shared HPC infrastructure so that they may derive the full benefit from its existence.

### The university administration

The university administration, with the CIO as the representative, is a crucial partner in the Contract for shared HPC infrastructure. Without the university administration as a partner at least at the minimal commitment specified below, the Contract cannot reach its common goals and will not work.

**Benefits.** By committing to a partner in the Contract for shared HPC infrastructure, the university obtains the following benefits.

1. The university derives significant public relations benefit from operating a shared HPC infrastructure as opposed to supporting individual faculty PIs and academic units to build numerous smaller facilities. The shared HPC infrastructure can have the size needed to be visible at the national level and enable UF to compete in the HPC arena.

---

<sup>2</sup> The Contract does not prohibit this, as that would be inconsistent with the mission of the university to support advanced research. There are cases where it is appropriate and advised for a faculty PI to build his or her own lab. However, the cost to renovate space for HPC is high and requires experience to be readily available in the unit: the cost to provide cooling includes labor and equipment costs and engineering and impact fees. Also hiring local support staff has the disadvantage of placing individuals in isolated situations with little prospect for raises and promotion. This condition often results in high turnover and lower quality of service for the faculty.

## HPC Contract

2. By operating a shared HPC infrastructure, it is possible to attract and retain highly talented individuals to staff the facility for system administration and for supporting researchers. Multiple smaller facilities cannot create an environment that is conducive to retention of such highly qualified and talented people.
3. By operating a shared HPC infrastructure, certain measures of standardization and uniformity can be put in place without compromising the mission of supporting advanced HPC research. The result are greater efficiency, stability and savings in cost.
4. A shared HPC infrastructure available to all faculty members on campus fosters interdisciplinary research by giving innovative faculty with an interest in HPC in a department that does not traditionally use HPC immediate access to a top quality facility and fosters communication with other faculty members with interest in HPC.

**Responsibilities.** The university commits to the following to support the operation and maintain the stability of the shared HPC infrastructure.

1. The university administration supervises the shared HPC infrastructure and ensures that appropriate governance structure is in place and functioning.
2. The university administration reviews the productivity of the shared HPC infrastructure and the effort and success rate of faculty members in submitting proposals and obtaining funding for research using HPC. It adjusts its annual stability commitment accordingly.
3. The university administration will provide negotiable matching funds for investments made by faculty PIs and academic administrators to ensure that the shared HPC facility is always better than the sum of the parts contributed by PIs and academic units.
4. The university administration will actively fund efforts to represent the university's HPC activity to the national and international community. This is in addition to the individual faculty PIs presenting their work at conferences and meetings.
5. The university administration also coordinates and fosters relationships with vendors regarding the shared HPC infrastructure.

## The HPC Center

The HPC Center staff members are the executive agents of the shared HPC infrastructure. They report administratively to the CIO. The HPC Committee,- and thus to the faculty and the academic administration,- acts as the steering committee for the HPC Center and determines what they work on and sets their priorities. The HPC Committee operates under the auspices of the Information Technology Advisory Committee (ITAC), which is advisory to the CIO.

**Responsibilities.** The HPC Center will implement the following in coordination with the CIO and the HPC Committee.

1. Report on the daily activity of the shared infrastructure, including usage statistics, maintenance operations, software and hardware upgrades.

## HPC Contract

2. Provide technical advice to the partners of the Contract for planning upgrades and acquisitions.
3. Interact with vendors and maintain vendor relationships.
4. Provide support for users in the form of online documentation, online issue tracking,<sup>3</sup> assistance in installation of software and training sessions.
5. Monitor the scheduling of jobs and ensure that fair access is provided to all stakeholders as specified in this Contract.<sup>4</sup>

---

<sup>3</sup> The HPC Center uses a web tool: <http://bugzilla.hpc.ufl.edu>.

<sup>4</sup> This is currently not quite possible, but the HPC Center is actively exploring several options to implement the scheduling capabilities required by the Contract.