

## University of Florida Campus Cyber-Infrastructure Plan 2011-2015

### Executive summary

As part of the strategy to make the University of Florida a member of the top 10 public universities in the United States, it is critical to build the right foundation for faculty and students to do their work in education and research. The University of Florida plans to build upon existing infrastructure and enhance it to reach the following goals and milestones:

#### - **Infrastructure**

- The 25,000 sq. ft. data center was completed in January 2013. It provides 5,000 sq. ft. of space for research computing and 5,000 sq. ft. for general enterprise computing services, including teaching and enterprise application support.
- Upgrade of the existing 20 Gb/s Campus Research Network to 200 Gb/s was completed on January 31, 2013 (with funding by an NSF MRI award). At the same time the connection to Florida Lambda Rail (FLR) was upgraded from 10 Gb/s to 100 Gb/s and the FLR link to Jacksonville from 10 Gb/s to 100 Gb/s (with partial funding by an NSF CC-NIE award).
- Join the Internet2 Innovation Platform<sup>1</sup> leveraging its 2004 investment in a scientific DMZ and work by its faculty on software defined networking (SDN).
- Added the BullGator supercomputer with 16,000 new cores to the HPC Center compute cluster for research computing to bring the total core count over 23,000 in March 2013. At the same time grow the research storage capacity to 5 PB, from 1.5 PB.
- Annually grow this compute and storage capacity to keep pace with active research projects. By 2015, the University of Florida will offer its faculty a balanced portfolio of compute and storage options ranging from on-campus, to cloud-provided, to national resources like XSEDE, all supported seamlessly by on-campus experts as well as online resources.
- Obtained InCommon Silver certification during 2012.

- #### - **Services and expertise**
- Build and continually grow a set of well-defined compute and storage offerings with expert consulting services and training to make researchers more efficient, productive, and competitive in obtaining funding for their proposed research projects. Establish a common strategy for BIGDATA across campus in collaboration between Research Computing and the libraries.

---

<sup>1</sup> <http://internet2.edu/news/pr/2012.04.25.disruptive-tech-to-advance-us-scientific-research.html>

- **Collaborative framework** With other universities in the State of Florida that are part of SSERCA, build and grow a cyber-framework consisting of compute and storage capabilities to support collaborative research projects across the state.

## Background

In April 2011 the University of Florida created coherent organization to support the computing needs of researchers. The organization is called UF Research Computing<sup>2</sup> and reports to the Office of IT with major additional funding from the Provost and the VP for Research. The new organization will build on the success of the UF HPC Center that was established in 2005 as a collaborative effort between faculty, departments, colleges and the Office of IT. UF Research Computing is governed by the Research Computing Advisory Committee<sup>3</sup> which sets strategic direction and policy. The committee is actively reaching out to include the digital humanities and social sciences.

An NSF MRI grant in 2004 established the Campus Research Network (CRN) as a separate 20 Gb/s network that connects several large-scale computing and storage facilities on campus and provides a 10 Gb/s link to the Florida Lambda Rail (FLR). This type of dedicated network for research is now called Scientific DMZ<sup>4</sup>.

## Goals for campus infrastructure

All scientific, engineering, scholarly, educational activities are deeply impacted by developments in information technologies. These innovations improve upon established approaches as well as make it feasible to follow previously unexplored paths to manipulate and investigate 'big data' in multiple ways. To ensure that its researchers, scholars and students remain productive and competitive in the coming decade, a university must provide the necessary framework and tools. There is a difference between knowing what these new tools can do, and physically knowing how to manipulate them in research. To help the first aspect, requires awareness-building activities, whereas the second is best supported by providing the tools in a way that is easy and intuitive to use. The evidence is in the fact that that is the standard set by many industry leaders such as Apple, Google, Amazon, Facebook, to name only a few, where anyone can quickly use new tools to perform complex transactions with little or no training.

The UF administration agrees that the time has come for the University to provide its faculty and students with computational facilities that are functional, cost-effective, and competitive with those of its peer institutions and is strongly committed to doing so. These facilities and the infrastructure supporting them will be built and maintained by

---

<sup>2</sup> Website <http://www.hpc.ufl.edu>

<sup>3</sup> Website <http://www.it.ufl.edu/governance/advisorycommittees/researchcomputing.html>

<sup>4</sup> <http://www.internet2.edu/pubs/Internet2-Innovation-Platform-FAQ.pdf>

scholar-professionals according to the strategic goals set forth by the IT governance structure, which consists mostly of faculty and students.

The foundation of a strong cyber environment is a coherent collection of well-managed resources. Although used in different ways and to different degrees across the various disciplines of science, engineering and humanities, we find that these resources primarily consist of the following:

- compute capacity and capability,
- storage capacity with provision for protection both from disaster and from unauthorized access,
- network capacity to create and share data sets large and small, and
- access to expertise for advice and training on how to use these resources.

It is crucial for a successful deployment of cyber infrastructure that the organization providing the resources is customer needs and service focused. Furthermore, by the very nature of research the needs cannot always be clearly specified with much advance notice. Therefore the service provider must be willing and able to support new and emerging scholarly needs in different disciplines. UF Research Computing is committed to meeting these requirements.

## Infrastructure

### Data center

The construction of the Eastside Campus Data Center started in November of 2011 and was completed date on January 1, 2013. The data center provides 10,000 square feet of machine room space and 2.25 MW of power. Half of the machine room space is dedicated to research computing, the other half is for enterprise computing, which includes campus web, email, and teaching services.

### Network capacity

An increasing number of research and education projects require fast access to large shared data sets. The University of Florida is connected to the Florida Lambda Rail at 10 Gb/s<sup>5</sup> which is being used at 30% capacity with peaks at 90%. A number of projects cannot be competitively developed and proposed without addressing how this limit will not adversely impact the project.

For this reason it is crucial for the University to invest in upgrading this capacity to 100 Gb/s and thereby provide its faculty with the environment in which they can properly develop innovative and transformative projects, for example, in engineering, life sciences and health research. The goal is to provide this upgrade during the academic

---

<sup>5</sup> Gigabit per second

year 2012-2013. This goal was reached on January 31, 2013 with partial funding from a NSF CC-NIE and MRI awards in 2012.

### Compute and storage capacity

Many projects require storage of large amounts of data and the concomitant compute capacity to process the data. It is increasingly difficult for individual research groups and departments to maintain the hardware systems to support such projects. The complexity of the systems, the requirements for high electrical power and cooling capacity, and the need for sophisticated security procedures make it more effective for the University to invest in professional staff to manage such systems for researchers and their collaborators. In this way, UF faculty remain free to focus on teaching and research rather than designing and operating complex data systems.

Within the academic year of 2012-2013, the resources operated by the UF HPC Center will add 16,000 cores<sup>6</sup> and 2.88 PB of storage. The acquisition will be partially funded by UF Research Computing and by principal investigators using funds from their projects augmented by the Research Computing Matching Program<sup>7</sup>. In addition, several clusters with GPUs will be added for highly parallel programming using this rapidly evolving technology. This goal was reached in March 2013.

The University will sustain and grow the compute, storage and network capacity to meet the demand from research projects. The funding model for the growth is based on a basic funding by the University augmented by funds from grants to individual projects. The computational capacity will be sustained between 20,000 and 30,000 cores, with demand for computational capacity beyond that being met by national centers and cloud service providers.

The storage capacity is likely to keep growing. However, there is a distinction that must be made between storage for temporary data processing and storage for permanent data archiving and sharing. The ratio of “working storage” to “archival storage” will be measured and adjusted to reach a sustainable equilibrium. The storage must also support collaboration within as well as beyond the campus.

---

<sup>6</sup> The system currently at position 100 on the November 2011 release of the Top500 list of super computers has 8,320 cores with a linear algebra benchmark performance of 126 Tflops. Adding 4,000 cores to the existing 6,000 cores operated by the HPC Center will raise the estimated HPL benchmark performance from 46 Tflops to 82 Tflops. This brings the University of Florida HPC cluster around position 170 in the November 2011 release of the top500 list of super computers that can be found at <http://top500.org/list/2011/11/100>. 1 Tflops stands for 1 Tera floating point operations per second; 1 Tera equals 1 trillion. On average the HPL benchmark performance for systems on the top500 list is equal to 75% of the theoretical peak performance. Real workload sustained performance varies widely and may be as low as 10% of peak. A measure being used by the US Government to assess useful performance is 30% of peak.

<sup>7</sup> The Research Computing Matching Program was started in the summer of 2011 and is again available in 2012. It is described in the document at <http://www.it.ufl.edu/governance/advisorycommittees/documents/research-computing-matching-program20120501.pdf>.

### **Integration with statewide infrastructure**

As collaborative research becomes the norm rather than the exception, it is imperative that the UF invests in supporting such projects with often complex and extensive requirements to enable collaboration. The University is a founding and active member of the Sunshine State Education and Research Alliance (SSERCA)<sup>8</sup>. During 2011 the organization has been defined by its founding members Florida State University, University of Central Florida, UF, University of Miami, and University of South Florida and bylaws have been developed. The alliance holds quarterly meetings hosted in turn by one of the member institutions.

Within the context of SSERCA, the University of Florida in cooperation with the other members, will develop a set of statewide services to provide compute and storage capacity that is easily accessible and shared by researchers working at any university in the state of Florida. The goal is to have an initial set of well-defined offerings by the end of 2013.

The compute services will be accessible using Condor as the common front-end to submit jobs, in addition to the Open Science Grid protocols for those researchers who are familiar with its use. The shared storage capacity will be made available at all participating institutions as a mounted file system on the respective HPC clusters operated by each member institution. Thus researchers from any SSERCA member institution can easily share large data sets for local and remote processing.

### **Implementation of IPv6 and InCommon**

The infrastructure to support IPv6 protocols has been planned and critical services such as routing and name service (DNS) have been implemented on test systems and networks. The core network of the campus network is capable of running IPv6. The current stage of the deployment and transition plan is engagement with the different units on campus to determine the details of how all colleges and departments will move their services from IPv4 to IPv6. The recent developments of moving more services and servers to central locations and the deployment of more private IP networks has brought some relief in the pressure on public IP addresses, making the transition less urgent. A process of revisiting network architectures to make sure that only those servers and services are on public IP addresses that really need them is ongoing and helps the preparation for moving to IPv6.

The primary system for campus-wide authentication is the GatorLink ID issued to UF students, faculty, staff, as well as alumni and various associated. The GatorLink ID is based on Shibboleth. During 2012, UF has achieved InCommon Silver certification.

---

<sup>8</sup> Website <http://www.sserca.org>