

HPC Annual Report 2008-2009

Sep 12, 2009

A new era for HPC at UF

With Phase III coming online on January 1, 2009, the strategic vision of the HPC Committee formulated in 2003 has been completed successfully. The University of Florida now has a fully functional Center for High-Performance Computing. President Machen has tasked the HPC Center and the HPC Committee to formulate a strategy for the next five years of the HPC Center.

Challenges ahead

The funding of the UF HPC Center remains a challenge that requires constant attention. In the current climate of shrinking budgets, this is especially challenging.

Staff At the end of the fiscal year 2008-2009, an agreement was reached that will provide stable funding for the three staff members of the HPC Center for the next three years, from July 1, 2009 until June 2012.

Facilities The Vice President for Research has provided significant funding to provide the facilities that house the Phase III cluster. The university will construct a building on the Eastside campus that will consolidate server room space for both administrative computing and HPC.

Hardware insertion To remain useful and provide competitive resources for the researchers at UF, the HPC hardware, consisting of compute nodes, network switches, and storage systems, must be renewed every three to four years. Until now the HPC Center hardware has been growing by adding about 1,000 cores every year. The first set of cores are now over 4 years old and need to be replaced in 2010 at an estimated cost of \$500,000 to \$750,000. However, the funds are not available. The hardware insertion is the part that, in the funding model used by the HPC Center, must come from faculty investments, e.g. paid with funds from grants. Faculty must support the HPC center by ensuring that research grants have a budget for computing. The current flow of such funds is not sufficient to provide a regular hardware insertion schedule. The University and College administration have made strong commitments to the long-term stability of the HPC Center.

To more effectively communicate to the faculty how the HPC Center operates, we plan the following activities throughout 2009-2010:

1. **Website** We will redesign the HPC Center web site to make information more easily available and more clear and compact;
2. **Service offerings** We will unveil a new slate of offerings so that more faculty will find that the HPC Center meets their HPC needs;
3. **Investment campaign** We will make short presentations to faculty meetings in relevant departments to ensure that faculty are aware of the need for funding and how to funds the HPC Center.

Coordination on the correct understanding what the HPC Center offers among Deans, Chairs and hiring committees is essential to ensure that new faculty members get the correct information.

Accomplishments and activities

Phase III The new cluster became available for production use on January 1, 2009. We had an official reception celebrating the event on April 30, 2009. Reprinted from IT Connections, the UF publication for IT staff.

ICBR and IFAS lead HPC Phase III Expansion

Photos: Tyler Jones, IFAS
Text: Erik Deumens, HPCC

In January 2009, the Phase III expansion of the UF High Performance Computing (HPC) cluster was completed. The effort to build the coalition started in early 2007, with planning complete in November of 2007. This expansion brought the total computing capacity of the HPC cluster to 2,500 cores, which translates to about 10 Tflops (floating point operations per second).

The Phase III expansion effort was spearheaded with crucial initial fundraising by the Interdisciplinary Center for Biotechnology Research (ICBR) and the Institute for Food and Agricultural Sciences (IFAS). The Large Hadron Collider (LHC) and collaboration and the Open Science Grid (OSG) collaborated and several faculty in CLAS also made direct investments.

IFAS Dean of Research, Mark McLellan, led the effort to bring HPC to the researchers in IFAS and committed the resources to join the HPC center in the Fall of 2007. Researchers in IFAS immediately started using the 1,600-core Phase II cluster for their projects. The HPC Center staff provided support that dramatically improved their productivity by using the cores in the cluster in parallel versus using them as individual compute servers. Since January 2009, IFAS computations have been running on the Phase III equipment. The majority of the IFAS computational work is composed of thousands of relatively short jobs. On a large, shared cluster like the 2,500-core HPC cluster, many cores may begin processing a large collection of jobs simultaneously so that they all are completed in parallel. Consequently the total time researchers must wait becomes very short. A smaller cluster with fewer cores could not process the same set of jobs as quickly as our current HPC configuration.

Working closely with ICBR and others, IFAS faculty recognize the HPC cluster as a way to extend our exploration of basic biology into a more comprehensive understanding of cellular function. Our goal is to explore everything from basic genomics to related

phenotype and map gene function to biological process. Our plant scientists are exploring key metabolic pathways and their connectivity to behavior affecting traits such as flavor & aroma volatiles and drought & salt tolerance. With a vision to the future, our agricultural engineers will model complex interactions between agricultural land uses and climate as well as adaptive responses to climate change. And from a global vision to a nano focus, we expect to map how nano-scale interactions might affect nutrient flow and transport dynamics in a soil matrix. Our world is full of compelling questions and the HPC is another important tool that will help us explore possible solutions. IFAS is proud to be a committed supporter of the HPC cluster.



IFAS Dean of Research McClellan (left) and ABE Chair Haman (right) listen as HPCC engineer Prescott (middle) shows the Phase III cluster nodes.

The Office of IT (OIT) helped to coordinate a small reception on April 30 to mark the Phase III expansion and rollout of a new computational software capability offered by ICBR that runs on the HPC Phase III hardware. President Machen and Senior Vice-President for Research Philips remarked on the crucial importance of high performance computing in modern research and expressed their strong support for the campus-wide HPC Center as a key example of faculty-driven effort. Prof. Avery commented on the value the HPC cluster brings to the OSG through the fast data transport connection provided by the 20 Gbps Campus Research Network, and the Florida Lambda Rail. Selvi Kadirvel, Electrical and Computer Engineering (ECE) graduate student with Prof. Fortes, demonstrated a software interface that greatly enables biological computing for BLAST on the HPC cluster. .



Six people in the front from left to right: President Machen, Senior Vice President for Research Philips, Prof McCarty, ICBR Assistant Director Nicklin, ICBR Director Ferl, IFAS-ABE Chair Haman.

The Phase III expansion is housed in a newly prepared machine room in 121Larsen Hall, home of the Department of ECE. The room was converted from an electrical engineering teaching lab in 2006 with funds provided by the Department of ECE, Dean Khargonekar of the COE and Dr. Hoit, Interim CIO. The machine room was equipped with power and cooling with substantial support from Senior Vice President for Research Philips in 2008 to support the current Phase III cluster with additional capacity for further expansion.

The HPC Center began at the initiative of Dr. Frazier Interim CIO. The HPC Committee he created recommended a three phase plan to build a HPC facility that serves all researchers on campus. Phase I was funded by the world-wide LHC collaboration and the OSG, both represented on campus by Prof. Avery in Physics, with matching funds from the dean of the College of Liberal Arts and Sciences (CLAS) and from the Office of IT. Phase I became operational in 2004. Phase II was funded by a number of faculty and departments in the College of Engineering (COE) with matching funds from the dean of COE and OIT. Phase II became operational in two stages in January 2006 and January 2007 with a total of 1,600 cores and 32 TB of fast storage. The LHC and OSG collaboration as well as several CLAS faculty members also invested in the second stage of Phase II.

Fundraising for Phase IV has begun and we invite any researcher who has a need for HPC IV, to contact the HPC Center for further information, and we welcome broad

participation in this vital faculty-led initiative. Requests to use HPC III can be made to HPC directly or through ICBR.

Funding agreement with CLAS, COE, IFAS, OIT A three-year agreement was signed, becoming effective July 1, 2009, with the College of Liberal Arts and Sciences, College of Engineering, Institute for Food and Agricultural Sciences, and the Office of Information Technology to support the salaries of the HPC Center staff for three years. Interim CIO Frazier completed the agreement which will allow the HPC Center to continue to operate with three highly-qualified staff members.

Budget overview

The HPC Center operates on a very small annual budget for operations, but it manages a significant amount of assets and funds. A brief overview is presented in the table below.

In the first part of the table, each row represents the entity that contributed funds towards the HPC Center operation. Faculty members contribute funds to buy equipment, individually or in small groups. Departments and Colleges and the Office of IT contribute funds for staff and facilities.

Description	2003-2004	2004-2005	2005-2006	2006-2007	2007-2008	2008-2009
Faculty	\$200,000	\$1,092,000	\$0	\$620,000	\$480,000	\$113,000
Dept/Inst	\$0	\$213,000	\$25,000	\$67,000	\$389,000	\$25,000
College	\$100,000	\$463,000	\$130,000	\$182,000	\$238,000	\$169,000
OIT	\$370,000	\$301,000	\$110,000	\$336,000	\$607,000	\$137,000
TOTAL IN	\$670,000	\$2,069,000	\$265,000	\$1,205,000	\$1,714,000	\$444,000
Staff	\$0	\$171,000	\$244,000	\$263,000	\$300,000	\$306,000
Facilities	\$70,000	\$0	\$0	\$158,000	\$340,000	\$0
Equipment	\$600,000	\$1,667,000	\$0	\$760,000	\$872,000	\$0
Operation	\$0	\$10,000	\$10,000	\$61,000	\$45,000	\$46,000
TOTAL OUT	\$670,000	\$1,848,000	\$254,000	\$1,242,000	\$1,557,000	\$352,000

List of investors

Some investors are investing as head of a college or department, some as individual faculty, some are playing both roles.

Archer D. L., Associate Dean, IFAS
Asthagiri A., Professor, Chemical Engineering
Avery P., Professor, Physics
Balachandar S., Chair, Mechanical and Aerospace Engineering
Beck H. W., Professor, Agricultural and Biological Engineering
Cheng H.-P., Professor, Physics QTP
Curtis J., Chair, Chemical Engineering
D'Anieri P., Dean, College of Liberal Arts and Sciences
Emond J.P., Professor, Agricultural and Biological Engineering
Frazier C., Interim CIO, Office of Information Technology
Gao B., Professor, Agricultural and Biological Engineering
Graham W., Professor, Agricultural and Biological Engineering
Fortes J., Professor, Electrical and Computer Engineering
Haman D.Z., Chair, Agricultural and Biological Engineering
Haselbacher A., Professor, Mechanical and Aerospace Engineering
Hoit M., Interim CIO, Office of Information Technology
Jones J., Professor, Agricultural and Biological Engineering
Khargonekar P., Dean, College of Engineering
Kiker G. A., Professor, Agricultural and Biological Engineering
Law M., Chair, Electrical and Computer Engineering
Martinez C., Professor, Agricultural and Biological Engineering
McLellan M., Dean of Research, IFAS
Merz K., Professor, Chemistry QTP
Munoz-Carpena R., Professor, Agricultural and Biological Engineering
Philpott S., Professor, Materials Science and Engineering
Shea J., Professor, Electrical and Computer Engineering
Sinnott S., Professor, Materials Science and Engineering
Soltis P., Professor, Florida Museum of Natural History

An additional group of investors consists of the principal investigators of the NSF-MRI 2004 grant that created the Campus Research Network and provided the InfiniBand switch network for the Phase II cluster with a matching grant from CISCO.

Ranka S., Professor, Computer Science and Engineering
George A. D., Professor, Electrical and Computer Engineering
Avery P., Professor, Physics
Trickey S. B., Professor, Physics
Sheng P., Professor, Coastal Engineering

Staff

The HPC Center reports to the office of the CIO, similar to CNS and AT, and has three permanent staff members:

1. Charles Taylor, IT senior, and associate director of the HPC Center,
2. Craig Prescott, IT senior,
3. Jon Akers, IT engineer

The part time position of program director is currently held by Dr. Erik Deumens, scientist in the Department of Chemistry and in the Department of Physics. In addition, the center has had several students working part-time.

Usage statistics

The HPC Center website provides a wide range of information, including live data about system utilization: <http://www.hpc.ufl.edu/index.php?body=util>. These are organized into three groups:

1. Cluster status shows the activity on the running system.
2. Queue status shows the jobs in the system, either running, or waiting for available resources.
3. Past status shows accounting statistics of jobs that are still running or have completed.

There are several displays worth pointing out.

1. “CPU Usage Summary“ <http://www.hpc.ufl.edu/index.php?body=pbs/nodestate> shows in a single view each of the 1600 CPUs with a color code indicating whether the CPU is idle or busy on a serial job, or a parallel job with 2-8, 9-32, 33-128, or 129 or more processors.
2. “CPU Job Utilization” <http://www.hpc.ufl.edu/index.php?body=pbs/cpustat> summarizes the same information with a bar chart and a pie chart.
3. “General Queue Status”, “Torque Queue Status” and “Moab Queue Status” show a full list of all jobs waiting and running with details such as number of CPUs requested and time accumulated in the queue or time accumulated executing.
4. “Cluster Usage Statistics (last 7 days)” show the percentage of the cluster time used labeled by research group as a pie chart and as a table. You can also request this information for a different number of days than 7.
5. “Usage by College” and “Usage by Department” show the percentage labeled by College and Department for the last 30 days. You can request any number of days at the bottom of the screen.

The table with the total hours used during fiscal year 2008-2009 organized by college, department, and research group identified by the sponsoring faculty member is shown below. Sums per department, and per college are listed, as well as the percentage of the total used by each college. The UF HPC Center served over 14 million CPU hours to its users.

Usage summary July 2008 – June 2009

College	Department	Investigator	Jobs	Hours	Total Hours	Hour %
CLAS						
	Physics	Alan Dorsey	159	1960.520607		
		Arthur Hebard	133	9743.962253		
		Brent Nelson	186	0.182777777		
		Christopher Stanton	268	514.170828		
		David Tanner	28	855.6905518		
		Ho-Bun Chan	180	878.035553		
		James Fry	25	22.69833374		
		Kevin Ingersent	8406	194578.0446		
		Khandker Muttalib	85	191.0722202		
		Mark Meisel	10	0.008888889		
		Paul Avery	978953	2736136.772		
		Peter Hirschfeld	35169	100116.7439		
		Robert Buchler	797	5754.780096		
		Total	1024399		3050752.683	21.58%
	Statistics	George Casella	390	5115.713647		
		Ronald Randles	5	4.267499924		
		Rongling Wu	143	2471.326175		
		Song Wu	368	2401.140442		
		Total	906		9992.447764	0.07%
	Zoology	Ben Bolker	5790	18531.37186		
		David Evans	52	8659.012017		
		Edward Braun	7	1125.599497		
		Michael Miyamoto	70	1098.276961		
		Rebecca Kimball	7	1191.260864		
		Total	5926		30605.5212	0.22%
	Chemistry	George Christou	255	7702.340599		
		John Eyler	39	513.1577759		
		John Reynolds	6	1.254999998		
		Mike Scott	143	28086.97536		
		Nick Polfer	12307	67338.89231		
		Nicole Horenstein	73	1058.349998		
		William Dolbier	252	14036.64256		
		Total	13075		118737.6136	0.84%
	Astronomy	Eric Ford	201788	700211.9247		
		Jonathan Tan	68	4450.76226		
		Michael Barker	19	803.9283447		

Annual Report 08-09

	Total	201875		705466.6153	4.99%
Anthropology	Connie Mulligan	12264	81956.41177		
	Total	12264		81956.41177	0.58%
FLMNH	Matt Gizendanner	2	0.186388895		
	Total	2		0.186388895	0.00%
Botony	Total	0		0	0.00%
Wildlife	Total	0		0	0.00%
Math	John Klauder	2537	35570.89635		
	Total	2537		35570.89635	0.25%
QTP	Adrian Roitberg	5963	359826.3645		
	David Micha	1536	57421.44394		
	Erik Deumens	74151	761231.1054		
	Hai-Ping Cheng	9801	1169275.869		
	Jeffrey Krause	24	9148.229054		
	John Sabin	23838	110166.177		
	Ken Merz	15605	743346.7947		
	Nigel Richards	1755	140112.4341		
	Rodney Bartlett	792	4386.939234		
	So Hirata	13192	51573.98321		
	Total	146657		3406489.341	24.09%
Biology	Total	0		0	0.00%
COE					
Coastal	Andrew Kennedy	95	7740.504554		
	Arnoldo Valle-Levinson	1651	168731.736		
	Peter Sheng	3112	4853.481858		
	Total	4858		181325.7224	1.28%
ISE	Panos Pardalos	32	67.23111062		
	Timothy Middelkoop	3814	19519.56582		
	Total	3846		19586.79693	0.14%
CISE	Anand Rangarajan	1	0		
	Arunava Banerjee	27372	84812.56216		
	Baba Vemuri	24	0.009166666		
	Total	27397		84812.57133	0.60%
EES	John Sansalone	53	570.8933439		
	Total	53		570.8933439	0.00%
NRE	Samim Anghaie	226	332.0366516		

Annual Report 08-09

	Total	226		332.0366516	0.00%
HCS	Alan George	3	30.99555588		
	Total	3		30.99555588	0.00%
ECE	Hanifph Latchman	547	3631.435791		
	Jing Guo	3698	14611.13226		
	John Shea	31283	632563.0961		
	Jose Fortes	1539	9560.787689		
	Tao Li	274233	234851.0903		
	Total	311300		895217.5421	6.33%
MAE	Nagaraj Arakere	8	0.002222222		
	Norman Fitz-Coy	6	13.80166666		
	S. Balachandar	12498	1786620.237		
	Subrata Roy	740	2197.695025		
	Youping Chen	6120	572110.8499		
	Total	19372		2360942.586	16.70%
MSE	Brij Moudgil	244	7495.001203		
	John Mecholsky	235	33117.04804		
	Shirley Meng	2539	269003.9343		
	Simon Phillpot	7735	632003.9311		
	Susan Sinnott	6406	553041.2423		
	Total	17159		1494661.157	10.57%
CHME	Aravind Asthagiri	1935	359539.9763		
	Dmitry Kopelevich	140	55456.24495		
	Jennifer Curtis	2350	237921.7911		
	Richard Dickinson	1	0		
	Tim Anderson	747	7382.626305		
	Total	5173		660300.6387	4.67%
BME	Huabei Jiang	8562	186454.2018		
	Mingzhou Ding	214	676.9563786		
	Total	8776		187131.1582	1.32%
ACIS	Jose Fortes	55662	55289.20328		
	Total	55662		55289.20328	0.39%
IFAS					
IFAS	Daniel Hahn	22	99.90333462		
	Donald McCarty	2830	1362.899698		
	James Jones	305	11757.14844		
	Jasmeet Judge	3300222	345124.153		
	Jean-Pierre Emond	2	0.000833333		
	Karen Koch	7	0.002777778		
	Rafael Munoz-Carpena	2919	17353.7403		

Annual Report 08-09

	Wendy Graham	475	45083.19876		
	Total	3306782		420781.0471	2.98%
Entomology					
	Total	0		0	0.00%
USFWS					
	Total	0		0	0.00%
COM					
ICBR					
	Li Liu	38723	84706.86196		
	Total	38723		84706.86196	0.60%
ICHP					
	June Nogle	60	1389.715576		
	Total	60		1389.715576	0.01%
MSG					
	Total	0		0	0.00%
UFMDC					
	Total	0		0	0.00%
CSB					
	Art Edison	668	121791.1555		
	Total	668		121791.1555	0.86%
Pathology					
	David Ostrov	92	982.7960868		
	Total	92		982.7960868	0.01%
MDPHD					
	Stephen Hsu	131	8354.700932		
	Total	131		8354.700932	0.06%
Physiology					
	Peter Sayeski	135	29764.26018		
	Total	135		29764.26018	0.21%
PHPH					
	Zhou Yang	1	5.845277786		
	Total	1		5.845277786	0.00%
MGM					
	Lauren McIntyre	157	8686.958008		
	Total	157		8686.958008	0.06%
BMB					
	Total	0		0	0.00%
EHPR					
	Yueh-Yun Chi	69	754.078064		
	Total	69		754.078064	0.01%
Pharmacology					
	Total	0		0	0.00%
Whitney lab					
	Total	0		0	0.00%
OIT					
HPC					
	Charles Taylor	19330	79681.87087		

Annual Report 08-09

	Total	19330		79681.87087	0.56%
FIU					
	Total	0		0	0.00%
SNRE					
	Stephen Humphrey	10372	3149.268066		
	Total	10372		3149.268066	0.02%
OIT					
	Total	0		0	0.00%
CBA					
Economics					
	Total	0		0	0.00%
Grand Total		0		14139821.58	100.00%