HARVARD UNIVERSITY
Faculty of Arts and Sciences
Research Computing, Division of Science

rchelp@fas.harvard.edu

Research Computing Newsletter – first quarter 2011 [ISSUE ONE]

Introduction

2010 was another amazing year for our Research Computing group. Yet again have we seen dramatic growth in the number of systems and servers we support; improving our efficiency and still averaging an additional 250 kW demand for power each and every year. We are executing more science – both for a larger user base and at larger scale. The power and space needed to support our >13,000 cpu cluster and large multi-petabyte data stores is close to exhaustion while the search for advanced HPC data center space continues and is being supported at the highest levels within the University. Fortunately we are directly involved in the Massachusetts Green High Performance Computing Center (MGHPCC) and have been working towards designing a state of the art data center facility in western Massachusetts with our colleagues at Boston University, MIT, Northeastern and the University of Massachusetts. We are on track with this monumental undertaking and plan to open the facility in the later half of 2012.

Additional Capacity

2010 saw an additional 1,300+ QDR connected Intel Westmere computing cores added to our physical sciences community. We have also added in excess of 2,000 additional CPU for individual faculty who have purchased these systems in concert with the Research Computing organization. All of this additional compute is tied to our unified auth and storage systems, effectively managed as a single device with sharing available to all via a managed queuing system. We have also introduced significant storage capacity for a number of groups across the cluster.

Our ganglia monitoring system is maintained with each cluster addition either added by the core or by working with our faculty. You can follow progress at:

http://software.rc.fas.harvard.edu/ganglia

New Virtual Machine Environment

Legacy VMware systems have been shutdown by effectively migrating systems to our new KVM infrastructure. The cluster is constructed from a 300TB iSCSI GFS fail over storage array. The GFS platform hosts the kernel virtual machines and is scalable by volume as we need to add additional capacity. There are 96 processors available on the physical hosts in a redundant topology. This allows us to provide a greater virtual machine capacity for lower cost and focus future investment on growing and upgrading the underlying hardware infrastructure. Virtual machines are available to any research group who needs one or more dedicated systems for unique computing needs.

Infrastructure Enhancements

Our 1 Summer facility is currently undergoing an in-place migration from two dedicated 60 ton chiller units over to a 640 ton multi-stack system to provide greater fault tolerance and scalability and to avoid a repeat of our recent outage. Our remaining on campus machine room space is also being modified with extra ducting, cooling and networking performance enhancements to allow us to continue to optimize our use of existing capacity while we wait for the MGHPCC facility to come on line.

Departmental updates

HUH - Harvard University Herbaria (HUH) and research computing have signed a collaborative support agreement to provide enhanced research computing support to the HUH. IT staff in the Herbaria are joining research computing full time to work on new and exciting projects together.

Upcoming maintenance June 10th 2011

Power systems maintenance in our 60 Oxford Street data center will be carried out over the evening of the 10th June 2011. Research Computing are taking advantage of this window to carry out some much needed maintenance of the cluster and file systems. Plan your in-silico experiments ahead by expecting that RC services will be unavailable from 4pm Friday June 10th to 10am Monday June 12th 2011. We will also take this opportunity to upgrade our LSF scheduling software.

More information at:

http://rc.fas.harvard.edu/updates/downtime_June2011

Welcome new team members

Two new folks have joined the research computing family recently. Keith Herron and Eric Mattison. Keith is responsible for our new high performance computing kernel virtual machine infrastructure. Keith joined us from a position in FAS IT. Eric is building out our advanced scientific application platform including SPINAL and CLEAN having joined us from Harvard's Decision Science Laboratory.

Museum Collections Systems

New virtual machines and database servers have been installed, allowing staff to focus on migrating collections, this is now complete. Our next stage is to upgrade our infrastructure both software and hardware to support our consolidated environment.

NCF 2.0

We continue to design new virtual services for containerized high-performance computing, extending work from 2010 to utilize kernel virtual machine environments (KVM) to provide research computing capacity securely on shared physical infrastructure. The Neurocomputing facility is our first production customer of this design. Work included migration of a multiple terabyte storage array to new fault tolerant storage systems.

CfA – The Center for Astrophysics and RC are searching for a Sr. Research Computing Associate to be split between RC and the CfA to support the Institute for Theory and Computation. This split hybrid model has worked exceptionally well for the department of Earth and Planetary Science for over two years; we aim to have identical success with the CfA!

Mathematics – Significant progress has been made through careful collaboration to provide new faster more fault tolerant servers and services to the department.

Physics – continued growth is driving new departmental file system installations.

IQSS – Driven by large data requirements from the Astrophysics community, RC and IQSS are collaborating to install the Dataverse Network (DVN) on RC hardware at scale with support and guidance from the software architects within IQ.

Life Sciences – increased adoption of the RC "SPINAL" system for monitoring, scheduling and managing over 100 scientific instruments. We continue to help researchers analyze high-throughput sequencing data for RNA-Seq and other emerging methods.

Center for Nanoscale Systems – further co development of instrument management systems continue with shared virtual machines and infrastructure.

Chemistry and Chemical Biology – new electronic notebook systems are installed and being actively used by our community. This is further reducing cost and improving our efficiency.

More Lustre upgrades are coming

We plan to migrate to newer more reliable high performance computing storage over the coming months. We will sunset the legacy /n/scratch systems and upgrade and enhance systems to support newer more reliable versions. Our initial 250TB increase is complete, patching is being performed in stages to upgrade the client nodes.

Windows High Performance Cluster

RC has built a Windows High Performance Computing (MSHPC) environment for research. It currently consists of 128 cpu hosted on both physical

Email Migration (June 10th 2011)

Research computing and the IRIS team have been focused on sun setting legacy email systems for the past 12 months. We are nearly complete, only a few accounts remain on the legacy systems. We are aiming for June 10th as a final cut over date to coincide with our systems downtime. Once complete this project allows us to migrate the legacy Active Directory system in the sciences to better support the Universal Desktop deployment. help@fas will soon be one single central place for both desktop and email support for the division, with rchelp@fas being a similar place for all your HPC, large storage and research computing support requests!

Digital Access to a Sky Century

The DASCH project is now well underway via high speed networking from the CfA to the Odyssey computing environment, allowing scanning and acquisition and analysis of 100 years (500,000) of plates to be digitized and preserved.

Retirement of legacy storage arrays

RC have migrated over ½ a petabyte of storage this year to preserve and extend the life of the data at rest. This has permitted us to retire our legacy storage and tape backup infrastructure and free up much needed space, power and cooling. We would also like to thank all of the groups that have worked with us as we migrated the data we could not have managed it without your support!

Security

RC and the University Security group continue to improve the security posture of their services. Advanced application scanning and detection are helping to keep Research Computing assets safe, secure and available.

RC top tip: *modulesearch*

Remember the UNIX command ***module avail*** will let you find out which software is loaded on the Odyssey cluster. You can also use ***module load hpc/rc*** to load in our ***modulesearch*** command that allows for fine grained search of our software stack.

and virtualized hardware, connected to shared storage. Windows 2k8 server and Windows XP OS platforms host diverse applications, focused on proteomics and mass spectrometry analysis. We plan to implement the LSF queuing system, to more tightly couple MSHPC with the Odyssey cluster. With growing demand, we plan to further expand Windows virtual compute nodes using Linux KVM.

Authentication Services

Our new high speed VPN/Active Directory and RSA security environment went live on the 14th of Feb 2011. Using a resilient multi-site redundant design with new dedicated high speed hardware, this successful deployment represents months of effort to build, test, and provide as seamless as possible transition for our >2,000 accounts while reissuing soft tokens.

Contact us

Never hesitate to reach out to us if you have a question or comment:

email at rchelp@fas.harvard.edu for any request

or you can leave us a message at our automatic 24/7 voice message service @ 617 299 9724.

Both methods automatically create RT tickets that each member on the RC team follows.

http://rc.fas.harvard.edu/faq

is also a great starting point for questions relating to the cluster and HPC/RC resources

We also have a twitter feed @fasrc you can follow for updates, this feed is great during system outages or major software updates.