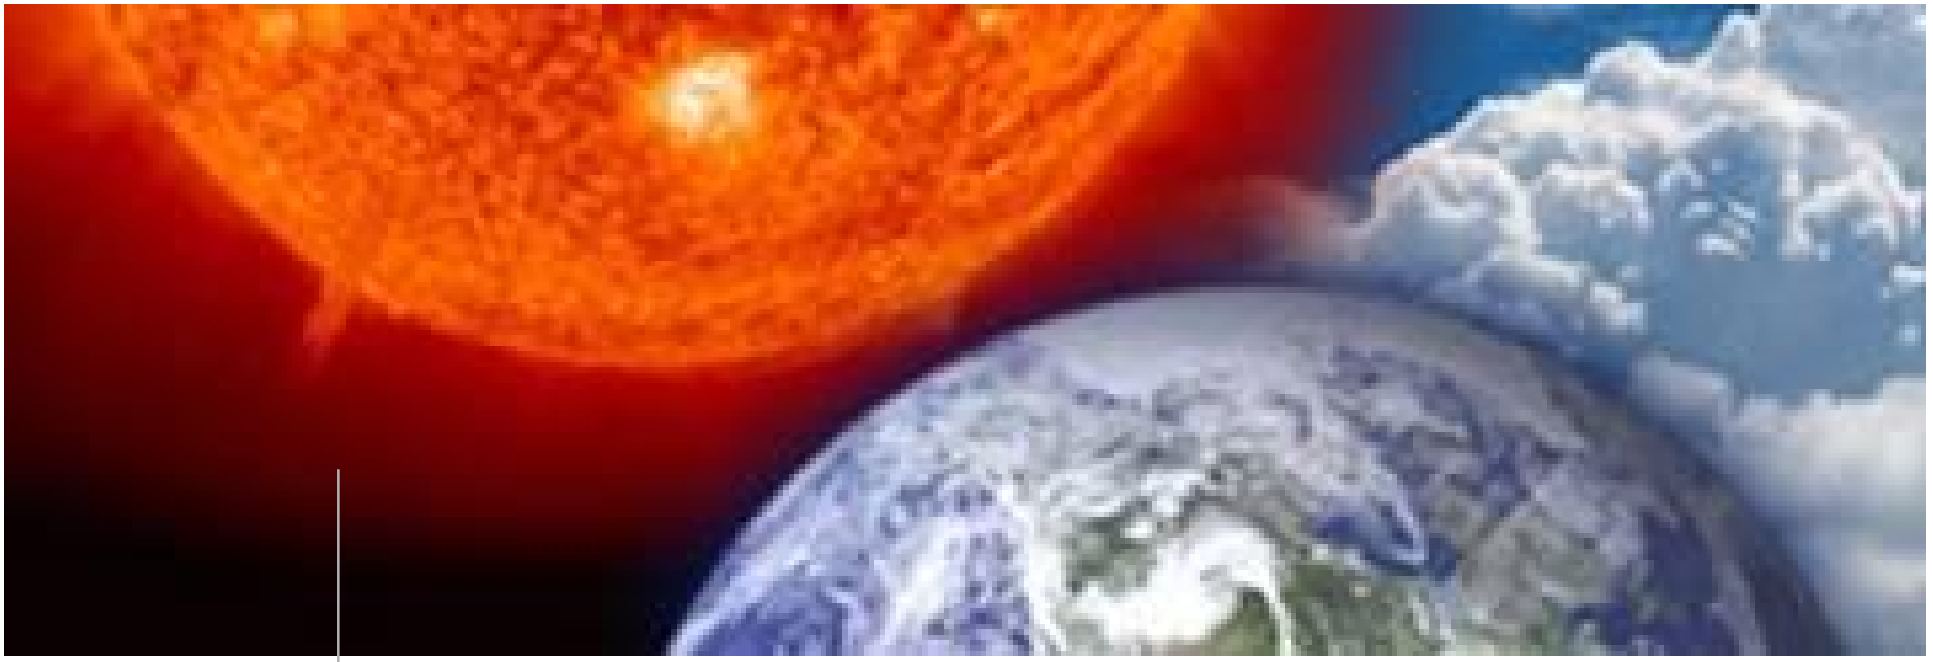

NCCS Cluster #2
Series of Brown Bag Presentations
October 2006

NASA Center for Computational Sciences (NCCS)
Computational & Information Sciences and Technology Office (CISTO)
Goddard Space Flight Center

Schedule over next few weeks

- 9 October
  - Final software installations, including PBS
  - Security scans, LDAP integration, etc.
- 16 October
  - Major power outage in install UPS
  - GPFS metadata configuration and testing
  - System on-line, very early access users
- 23 October
  - Finalize GPFS
  - Fail over configuration
  - Early access users
- 30 October
  - Pioneer users
- 6 November
  - Beginning of acceptance
  - General user access

Home, nobackup, project file systems

File System	Type	Login	Gateway	Compute
/discover/home/<username> • symlinked to /home/<username> ($HOME) • potentially make this available on explore • quotas with initially small storage space (~500 MB)	GPFS	Yes	Yes	Yes
/discover/nobackup[1-4]/<username> • symlinked to /nobackup/<username> ($NOBACKUP) • potentially make this available on explore • quotas with initially a large storage space (~100 GB) • not skulked initially (possible GPFS skulking)	GPFS	Yes	Yes	Yes
/explore/home/<username> • NFS mounted from explore	CXFS NFS	Yes	Yes	**No**
/explore/nobackup[1-4]/<username> • NFS mounted from explore	CXFS NFS	Yes	Yes	**No**
/discover/nobackup[1-4]/project/<projectname> • project specific directories • space driven by project requirements	GPFS	Yes	Yes	Yes

Other file systems

File System	Type	Login	Gateway	Compute
/usr/local • /usr/local/otherlibs will be created under here (or whatever you want to call it)	GPFS	Yes	Yes	Yes
/usr/local/unsupported • symlinked to /usr/ulocal	GPFS	Yes	Yes	Yes
/usr/nlocal • node specific local directory (NOT Global)	Local	Yes	Yes	Yes
/opt • MPI libraries stored here • Don't need to worry about this, just load the module • Same on each node (local to each node)	Local	Yes	Yes	Yes
/g[1-8] and other DMF archive directories • symlinked to /archive/<username> ($ARCHIVE) - Details TBD	CXFS NFS	Yes	Yes	**No**

Currently Available Modules

borgmg:/ # module avail

------------------------ /usr/share/modules/modulefiles ------------------------

bgcc/autoconf-2.13	**comp/gcc-3.4.6**	**module-info**
bgcc/gcc-3.2.3	**comp/gcc-4.0.3**	**modules**
bgcc/gcc-3.3.1	**comp/gcc-4.1.1**	**mpi/intel-2.0.1**
bgcc/gcc-3.3.3	**comp/intel-8.1.034**	**mpi/mpich-1.2.7**
bgcc/gcc-3.3.6	**comp/intel-8.1.038**	**mpi/openmpi-1.1**
bgcc/gcc-3.4.3	**comp/intel-9.1.038**	**mpi/scali-5.1.0.1**
bgcc/gcc-3.4.6	**comp/intel-9.1.039**	**mpi/sst-3.3.0.4.1-crc**
bgcc/gcc-4.0.3	**comp/intel-9.1.042**	**mpi/sst-3.3.0.4.1-crc2**
bgcc/gcc-4.1.1	**comp/intel-9.1.043**	**mpi/sst-3.3.0.4.1-crc4**
comp/gcc-3.2.3	**comp/pgi-6.1.6**	**mpi/sst-3.3.0.4.1-crc5**
comp/gcc-3.3.1	**dot**	**mpi/sst-3.3.0.4.1-crc6**
comp/gcc-3.3.3	**lib/mkl-8.1**	**mpi/sst-3.3.0.8.4**
comp/gcc-3.3.6	**lib/mkl-9.0beta**	**null**
comp/gcc-3.4.3	**module-cvs**	**use.own**

Many of which need not be available to the users.

User Available Modules
Something like the following...

borgmg:/ # module avail

---------- /usr/share/modules/modulefiles ----------

comp/gcc-3.4.6	**module-cvs**
comp/gcc-4.1.1	**module-info**
comp/intel-8.1.034	**modules**
comp/intel-8.1.038	**mpi/intel-2.0.1**
comp/intel-9.1.038	**mpi/mpich-1.2.7**
comp/intel-9.1.039	**mpi/openmpi-1.1**
comp/intel-9.1.042	**mpi/scali-5.1.0.1**
comp/intel-9.1.043	**mpi/sst-3.3.0.8.4**
comp/pgi-6.1.6	**mpi/sst-3.3.0.8.4-crc**
lib/mkl-8.1	**use.own**
lib/mkl-9.0beta	

Only keep a small number of gcc compiler versions, and a relatively manageable set of Intel compiler versions.

As versions are retired, permissions will be changed so they do not show up with a module available listing, but they will still be on the system.

Module Display

```
borgmg:/ # module display comp/intel-9.1.042
-------------------------------------------------------------------
/usr/share/modules/modulefiles/comp/intel-9.1.042:

conflict        comp
module-whatis   loads the comp/intel-9.1.042 environment
setenv          INTEL_LICENSE_FILE /usr/local/intel/license
setenv          COMPILER intel
setenv          COMPILER_VER 9.1.042
setenv          CC /usr/local/intel/cce/9.1.042/bin/icc
setenv          FC /usr/local/intel/fce/9.1.036/bin/ifort
prepend-path    PATH /usr/local/intel/cce/9.1.042/bin:/usr/local/intel/fce/9.1.036/bin
prepend-path    LD_LIBRARY_PATH /usr/local/intel/cce/9.1.042/lib:/usr/local/intel/fce/9.1.036/lib
prepend-path    LIBRARY_PATH /usr/local/intel/cce/9.1.042/lib:/usr/local/intel/fce/9.1.036/lib
prepend-path    INCLUDE /usr/local/intel/cce/9.1.042/lib:/usr/local/intel/fce/9.1.036/include
prepend-path    MANPATH /usr/local/intel/cce/9.1.042/man:/usr/local/intel/fce/9.1.036/man
-------------------------------------------------------------------
```

Moving Data - Gateways

- Specialized nodes for data movement
  - Two (2) gateway nodes for the base unit configured with the following
    - Dual socket, dual core Dempsey
    - 8 GB of RAM
    - 10 GbE network interfaces
- Scheduled via PBS specifically for data movement queues
  - Single CPU jobs (for data movement only)
  - “datamove” queue will run on these nodes

Visualization Nodes

- 16 nodes configure into the base unit
  - AMD Opteron dual core 280 processors (2.6 GHz)
  - 8 x 1 GB of PC3200/DDR400 S/R DIMM
  - 250 GB SATA
  - PCI-Express with SilverStorm inifinband 4x HCA (10 Gb)
  - NVidia Quadro FX 4500 PCI-Express
  - 10 GbE network interfaces
- User Environment
  - Same user environment as all other nodes, including additional software like IDL
  - Same file systems and tools
- How will this be used?
  - Future plans call for a potential viz wall
  - Post process, analysis, special processing

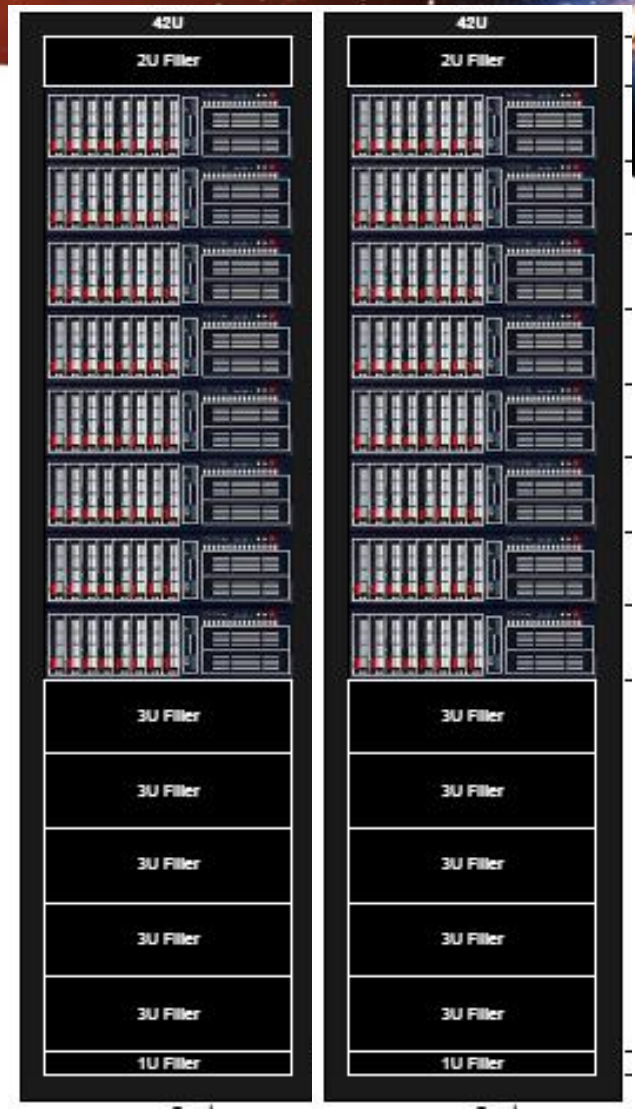

Test and development system

- A small test and development system is being configured to be as identical to the production system as possible
  - Same file system configuration with reduced number of data and metadata servers but with the same failover setup
  - Same management server configuration with failover setup
  - Reduced number of login and gateway systems, but can mimic necessary DNS and other failover scenarios
  - Similar disk subsystems, but not identical
- Is it big enough to test applications?
  - Configuring the system with 10 nodes of Dempsey and will expand to include 10 nodes of Woodcrest later
  - 40 cores of each processor: is that enough?
  - NOT a production system