

**NOAA R&D HPCS RFP
Appendix A.1**

~

**Current R&D HPCS Architecture
Boulder**

07 October 2009

Boulder Facility Organization and Subsystems Overview

The Boulder HPC systems are divided amongst two rooms in the David Skaggs Research Center (DSRC) in Boulder, CO. Room 2B407 is primarily for low power density equipment such as storage and HSMS. Room GA405 is for high density equipment such as compute racks, although some storage is still in GA405 and some compute is installed in 2B407. Both rooms are connected via 16 active, 150 meter, QDR IB links so that storage is visible across all computer systems.

There are currently 4 HPC subsystems at Boulder. The wJet, hJet and GPU systems are available for use by the general Boulder HPC users. The nJet system is dedicated to users of the Hurricane Forecast Improvement Project.

Subsystems Overview

System	Wjet	Hjet	GPU System	nJet
Installation Date	11/2006	8/2008	12/2008	8/2009
Node Count	364	252	16	448
Node Type	Intel Woodcrest	Intel Harpertown	Intel Harpertown	Intel Nehalem
CPU Speed	2.66 GHz	2.8 GHz	2.8 GHz	2.8 GHz
Cores/node	4	8	8	8
Total Cores	1440	2016	128	3584
Mem/core	3GB	2GB	2GB	3GB
Infiniband Speed	SDR	DDR	SDR	QDR

Note: The GPU System includes 2 NVIDIA GTX280 video cards for GPU processing.

Filesystems

Home File System (HFS)

- Netapp FAS 3140
- Redundant controllers
- 4 tiers of disk (144GB/15k rpm FC)
- 6 ethernet ports per controller
- Useable capacity approximately 6 TB

High-Performance Filesystems

- lfs0
 - DataDirect/Lustre filesystem
 - Single DDN9900 couplet
 - 4 OSS servers/
DDR Infiniband connected
Configured in pairs for failover
 - 2 MDS servers. Infiniband connected
Configured for failover
 - Usable capacity approximately 300TB
 - Measured performance approximately 5.5 GB/s
- lfs1
 - DataDirect/Lustre filesystem
 - Two DDN9900 couplets
 - 8 OSS servers
DDR Infiniband connected
Configured in pairs for failover
 - 2 MDS servers. Infiniband connected
Configured for failover
 - Usable capacity approximately 340TB
 - Measured performance approximately 10 GB/s
- panfs0
 - Panasas filesystem (To be installed October 2009)
 - 11 Shelves, each with 1TB disks
 - All shelves have their own controller
 - Includes productized Panasas Ethernet-to-IB gateway
 - Usable capacity approximately 170 TB
 - Estimated performance to be at least 5.5 GB/s
- lfs2 (Installation to be complete November 2009)
 - Vendor supported/open source Lustre implementation
 - 15 OSS servers
No failover between OSS servers
 - 2 MDS
Configured in failover pair
 - Usable capacity approximately 66 TB

HSMS

- Robot 1, AML/J
 - 6 LTO1 drives
 - 4 LTO3 drives
 - 1800 slots
 - To be decommissioned April, 2010

- Robot 2, Quantum I2000
 - 5 expansion units
 - 1 LTO3 drives
 - 5 LTO4 drives
 - 1800 licensed slots

- Disk cache (upgrade to be completed November 2009)

- DDN 9550
 - 500 TB SATA drives
 - Usable capacity approximately 51 TB

- Stornext 3.5