



Monash Central High Performance Computing (HPC) Facility



Monash Sun Grid (MSG)

The Monash Sun Grid (MSG) is Monash University's central production-class High Performance Computing (HPC) and Data Storage facility. First nodes were established in 2005. It has been expanded annually since, funded from central Monash operational and capital development funds. The MSG supports a wide variety of Monash researchers spread across most

faculties, with 80% continuous utilisation of the HPC service, and 1.3 petabytes of data storage (LaRDS).

Enquiries: merc@adm.monash.edu.au

2009 Configuration

The 2009 MCG HPC configuration provides 5 classes of HPC 'execute node' capacity:

- **General purpose (6.4 TeraFlops):**

640 CPU cores, each with 2 GB ram/core, plus dual 146 GB hard drives per node, which support virtual memory system to allow 4 GB peak memory per job 'slot'. Hardware comprises Sun X6250 blades, each with two quad core intel L5420 CPUs, total 640 cores at 10 Gflop/core (4 flop/Hz). Very low power consumption achieving >200 Mflop/W or ~25 W/core typical.

- **Hi-RAM (624 GFlops):**

20 CPU cores, each with 10 GB ram/core, plus dual 80~146 GB hard drives per node, which support a variety of flexible RAM config options, up to 48 GB per job slot upon request. Hardware comprises Sun X6220 blades, each with two dual core AMD 2218 CPUs. Total 120 cores at 5.2 Gflop/core.

- **Tightly-coupled (900 GFlops):**

96 cores, at 3 GB RAM/core, plus dual 146 GB hard drives per node. Hardware comprises Sun X4600 servers, each with 32 core shared memory architecture comprising eight quad core AMD 8356 CPUs and 96 GB RAM. High-bandwidth inter-process communication comprising: AMD hypertransport for very high bandwidth inter-process/shared memory communication within the 32 cores in each server, and infiniband between the 3 servers. Total 96 cores at 9.2 Gflop/core (4 flop/Hz). These are proving very popular already with MPI-intensive users (commissioned 2009).

- **Distributed Computing CONDOR Pool:**

>1000 cores spanning several faculties, campuses and divisions, including ITS labs. Participating faculties include: the Faculty of Medicine, the Faculty of Science and the Faculty of Information Technology. The Monash SPONGE uses Condor to execute Java, C, C++, Fortran-77 and Globus jobs. This Condor pool soaks up spare CPU cycles that would otherwise be wasted within operating desktop computers around the Monash campuses.

- **Interactive:**

For interactive login, visualization, animation and real-time results processing – refer to head-node facilities below for details.

Common Features

- Head-nodes: Dual Sun X4600 each with 16 cores (expandable) at ~4 GB/core and dual 146 GB hard drives per node, running VMware. Supports a variety of virtual machines to provide the following functions: login node; visualization node; interactive front-end; SGE job despatch; license server; system boot; Condor headnode; and various system management functions.

- Linux operating system.

- Dual hard-drives: These are configured in a RAID0 (mirror) configuration to allow system continuation in the even of disk failure – important for HPC jobs, which are typically long-running.

- Dual PSUs: Again to achieve high availability for long-running jobs.

- Dual gigabit ethernet connections to Monash network for: inter-node (MPI) communications; communications with head node machines; and communications with (LaRDS) storage system.

- **Summary:**

Total MCG execute capacity (2009 configuration):

920 cores; 2.8 TB RAM; 8 Tflop/s; 40 kW + >1000 Condor Cores.

Data Storage System

Each execute node, as well as the head nodes, of the Monash Sun Grid has direct gigabit ethernet NFS access to Monash's central research data store (LaRDS).

- The central LaRDS data store comprises dual SGI Altix A450 multi-CPU servers; 400 TB of disk; and dual IBM 3500 tape libraries having a current capacity of 1.7 PB (expandable as required). Servers and storage capacity are distributed equally across Monash's two main data-centres via Cisco MDS 9500 fibre-channel SAN. Storage capacity is managed using SGI's Data Migration Facility (DMF) hierarchical storage management (HSM) environment.

List of Current Software in Active Use on MSG

- **Molecular Dynamics code:**

NAMD, GROMACS, AMBER 10

- **Computational Chemistry**

Gaussian 03
GaussView 4
GAMESS
Molden

- **Computational Biology:**

Blast, clustalw, and related code

Fluent -- Computational fluid dynamics (licensed to Engineering)

XMDS -- ODE and PDE solver

MATLAB

R -- Statistical package

Ox -- Object-oriented matrix programming language

FSL -- Comprehensive library for analysis of FMRI, MRI and DTI brain imaging data

CircuitScope -- Open-source code for landscape genetics research and conservation planning using circuit theory

FFTW -- Fast fourier transform library

ENVIMET -- Micro climate model

ENVI/IDL -- Image processing system for environmental / geography research (running on front node)

- **Software Development tools, libraries, and parallel environments:**

OpenMPI

MPICH2

Python

NetCDF

PETsc

Java Compilers

Intel Fortran and C/C++ compilers

Intel Cluster Maths Kernel Library

(BLAS, LAPACK, etc.)

Contact MeRC

MeRC support staff are available to discuss all your research needs.

*Please email your general enquiry to:

merc@adm.monash.edu.au

- or phone:

Anthony Beitz, Technical Manager, MeRC
on 990 58681

Neil Clarke, Director of ITS Research
Shared Services on 990 54775

Wojtek Goscinski, Advanced
Technology Analyst, MeRC on 990
20715