




File Systems

Advance Science
with Pawsey Supercomputing Centre



Australian Government



File Systems

- File systems are storage areas mounted to the supercomputers and used by supercomputing jobs for data input and output (I/O).
- Three file systems are available to Pawsey users: **scratch**, **group**, **home**
- Intended for short-to-mid term storage
- **scratch** and **group** are not backed up, only **home** is
- For longer term storage, email help@pawsey.org.au to find out about RDS and HSM
- **Demo:** `df -h`
- **scratch** (`/scratch` and `/scratch2`), **group** (`/group`) are connected via Mellanox Infiniband
- **home** (`/home`) is connected via Ethernet
- There are no local disks to the nodes.

Shared across Pawsey

- **scratch, group** and **home** are global to
 - all nodes (login, compute, data mover)
 - and all Pawsey machines (Galaxy is a bit different)
- Convenient way of accessing the same files from different platforms

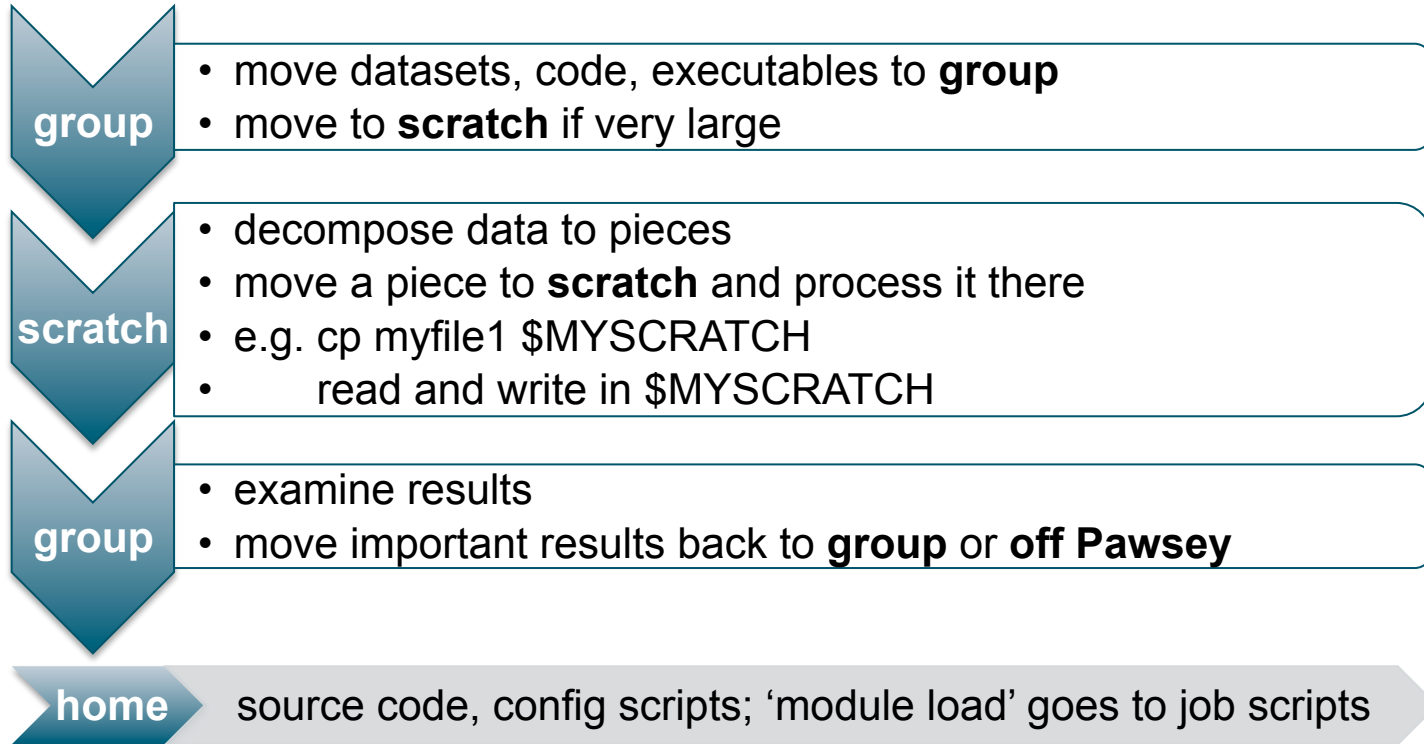
Machine	scratch	group	home
Magnus	/scratch	/group	/home
Zeus/Zythos	/scratch	/group	/home
Galaxy	/scratch2	/group	/home

File Systems

File System	Type	Features	Speed	Size	Quota	Purged	Back-up
/scratch	Lustre	Parallel IO	~72Gb/s	3PB	None	Yes, 30 days	No
/scratch2	Lustre	Parallel IO	~56Gb/s	1.5PB	None	No	No
/group	Lustre	Parallel IO	~30Gb/s	1.5PB (after expansion)	1TB per group	No	No
/home	NFS	-	~10Gb/s	15TB	10GB per user	No	Since Q2 2015

File System	Location	Permissions	Intended Uses	Lifetime
\$MYSCRATCH	/scratch/projectID/username /scratch2/projectID/username	750	Runtime IO: large reads/writes by supercomputing jobs	Subject to purge
\$MYGROUP	/group/projectID/username	750	Project-based storage: source codes, executables, datasets, job scripts, etc, shared within the group	Lifetime of project
\$HOME	/home/username	700	User-centric storage: login scripts, configuration scripts, etc	Lifetime of user account

Typical Workflow



Check Quota/Usage

- Both **scratch** and **group** are Lustre file systems.
Demo: `lfs quota -g courses01 /group` (no quota on scratch)
- **home** is NFS file system.
Demo: `quota`
- Alternatively, use pawseytools module.
Demo: `pawseyAccountBalance -project=courses01 -storage`



UNIX File Groups

- Used to control file permissions and manage quota/usage
- **Demo:** `ls -l myfile`
`id cou000`
`chgrp courses01 myfile`
- Every user is at least a member of two Unix groups: username + projectID
- File is associated with your default group when created.
- User can change associated group to other groups he/she is a member of.

File Permissions

- Access control
- **Demo:** `ls $MYGROUP/cou001; ls /home/cou001`
- A user's **home** is only accessible by that user.
- A user's **group** and **scratch** are accessible by that user and members of the same project.
- Share data within people from the same group
- (and not share with people who is not a member of the group)



Quota/Usage

- Manage quota/usage
- Files owned by **username** group count against 10GB on **home** and **group**
- Files owned by **projectID** group count against 1TB on **group**
- There is no quota on **scratch**
- Error: Disk quota exceeded!
- Files are owned by username group and hence counted towards the small quota
- **Demo:** `ls -l mydir ; chgrp -R courses01 mydir`
- change *mydir* and the files in it to be associated with *courses01* instead of *cou000*

Quota/Usage

- Group sticky attribute
- `cd $MYGROUP/mydir ; touch newfile`
- `newfile` is still associated with `cou000` instead of `courses01`
- `chmod -R g+s $MYGROUP/mydir ; touch newnewfile`
- Add group-sticky bit to `mydir` and `newnewfile` is now associated with `courses01`
- **Use Pawsey tool to fix group permissions!**
- `fix.group.permission.sh courses01` (part of 'pawseytools' module)



Strongly
Recommended!!



Thank You!

help@pawsey.org.au



Australian Government



- **User Documentation**
- <https://support.pawsey.org.au/documentation/>
- **User Training**
- <http://www.pawsey.org.au/training/>
- **Subscribe/Follow us**
- Pawsey Friends, Facebook, Twitter, YouTube
- **Acknowledge Pawsey**
- This work was supported by Pawsey Supercomputing Centre through the use of advanced computing resources.

