



Supercomputing User Training



Module 3: Filesystems Overview

Pawsey Training Series

Supercomputing User Training

1. Supercomputing Introduction
2. Logging In
3. Filesystems Overview
4. Moving Data In and Out
5. Using Software Modules
6. Using Software Containers
7. Accounting Model Overview
8. Job Scheduling Overview
9. Running Jobs
10. Testing Job Runs
11. Managing Project Data


Outcomes for this Module

- Describe what is a filesystem
 - List and describe the purposes of the various filesystems at Pawsey
 - Monitor your filesystem quotas using the shell
-
- ✓ Prerequisite knowledge [optional]:
 - ✓ **Bash shell basics**
 - ✓ **User Training 02: Logging In**

Watch for These Signs!



Definition of new concepts



Hands-on coding (demo)



Best practices



Exercises and solutions



Warnings (bad practices)



Links to user documentation

Available Filesystems at Pawsey



Australian Government



NCRIS
National Research
Infrastructure for Australia
An Australian Government Initiative



CSIRO



Curtin University



Murdoch
UNIVERSITY



GOVERNMENT OF
WESTERN AUSTRALIA



ECU
EDITH CURRIE



THE UNIVERSITY OF
WESTERN
AUSTRALIA

Why are we talking about filesystems anyway?



Filesystem

A method used by the operating system of a computer to manage, store and retrieve data.

In this context, we are also implicitly referring to the underlying hardware used to store data.

A personal computer typically comes with one type of filesystem and disk drives, for all uses

Examples of filesystem:

- Linux: “Ext4”
- Apple: “APFS”
- Windows: “NTFS”

Examples of hard-drives

- Spinning disks
- Solid state drives

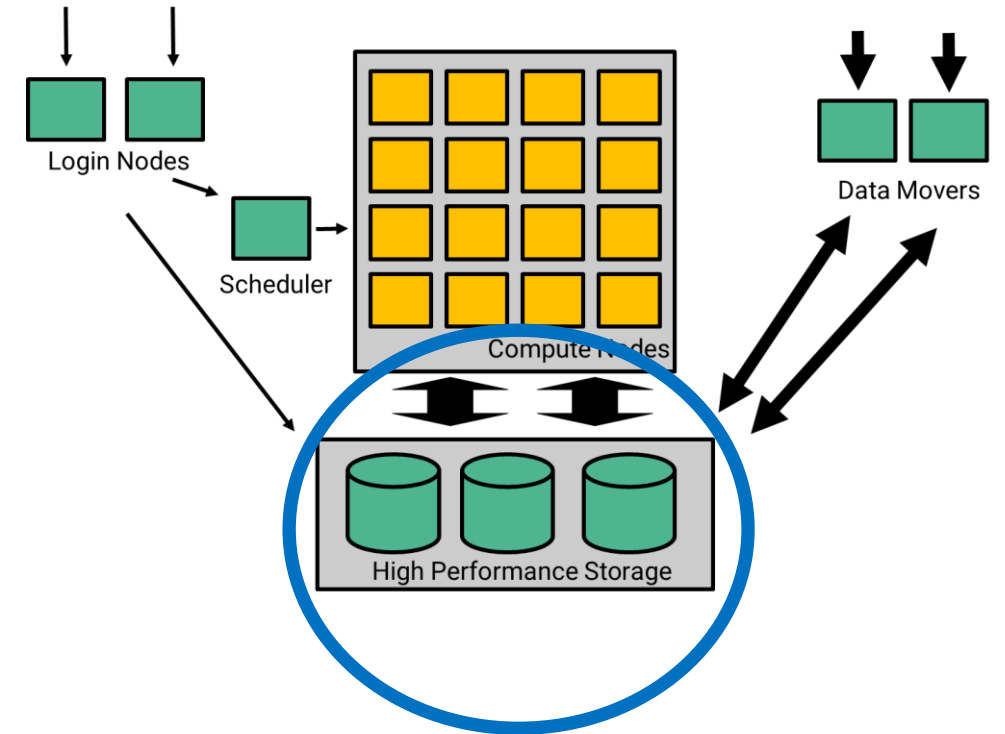


Why are we talking about filesystems anyway?

Filesystems in a Supercomputer:

To maximise performance, filesystems (and storage hardware) in supercomputers are specialised for various, distinct purposes

- Multiple types of filesystems available
- Specific policies are enforced on each type of filesystem
- Some optimised for read/write performance (fast but expensive)
- Some optimised for storage (power efficiency)
- Most are also optimised for shared use by multiple users at once
- Some not optimised at all (cheaper)



NOTE: in this module we are not discussing the specifics of storage hardware


Filesystems in Setonix

General Purpose


Working Data



Software



Project Data



Configuration Files



Dedicated to Radio-astronomy

- /astro
- /askapbuffer
- Mounted on the Garrawarla supercomputer
- Mounted on a dedicated data mover node (in the future)



More details @

- [Pawsey Filesystems and their Usage](#)
- [Filesystems Policies](#)



Working Data: The /scratch Filesystem

- Designed for fast, parallel, multi-user access of data
 - Technically a “Lustre parallel filesystem”
- Suitable as temporary work space for queued, running, or recently finished jobs
- Large 15 PB total capacity
- Quota limits:
 - Hard limit of 2PB per project and 2 million files per user
 - Soft limit of 1PB per project and 1 million file per user
- User directory is */scratch/project-id/username*
 - Available as **\$MYSCRATCH** environment variable
- Files on */scratch* are purged after 30 days



Be mindful of the 30-day purge policy on /scratch.

- Move out files from scratch you want to **store**
 - **Delete** files that you do not need any more
- This good practice ensures optimal performance for all users

Project Data: Acacia Object Storage

- Designed for high-availability, large-volume data storage
 - Technically “Ceph object storage”
- Suitable as storage space for the entire lifecycle of the project allocation
- Merit allocation supercomputing projects have 1 TB of project-wide allocation by default
 - Up to 10 TB can be requested via the Service Desk
 - Larger allocations must apply separately
- Provides different functionality and operates differently to traditional filesystems such as `/scratch`
 - Access using S3 interface
 - Data on object storage is read-only

The /software Filesystem

- Designed for fast, parallel, multi-user access of software
 - Technically a “Lustre parallel filesystem” (same as /scratch)
- Suitable for Pawsey- and user-managed software, workflow batch scripts and workflow configuration files
- Quota limits of 256 GB per project and 100,000 files per user
- Organised as:
 - Project-wide level: */software/projects/project-id*
 - User level: */software/projects/project-id/username*
 - Latter available as *\$MYSOFTWARE* environment variable
- Data retained for the lifetime of the project



The /home Filesystem

- Not designed for performance, nor for large-scale parallel access
 - Technically a “NFS filesystem”
- Suitable only for software configuration files, such as `.bashrc`, and user private data, such as access keys
- Small quota limits of 1 GB per user and 10,000 files per user
- User directory is `/home/username`
 - Available as the `~` symbol or the `$HOME` environment variable



Avoid keeping workflow-related settings inside shell configuration files.

Define workflow settings in Slurm batch scripts instead. Doing so improves the reproducibility of your workflows.

Monitoring Filesystem Usage



Demo on Setonix – let's do this together

- Filesystem usage can be checked from the command line:
 - For `/home`: `quota -s` command
 - For `/scratch` and `/software`: `pawseyAccountBalance -s` command
 - More details via the Lustre command `lfs quota -h <options>`



Proactively monitor both file size and file count

This practice will reduce your likelihood of hitting the quota limits; whenever this happens, no files can be written until usage is brought back below quota.

OUTPUTS: Monitoring Filesystem Usage

```
$ quota -s
Disk quotas for user username (uid userid):
  Filesystem  space  quota  limit  grace  files  quota  limit  grace
172.18.0.100:/home
                137M  1024M  1024M          312  10000  10000
```

```
$ pawseyAccountBalance -s
Compute Information
```

```
-----
Project ID      Allocation      Usage      % used
-----
pawsey0001      25000           18019      72.1
```

```
Storage Information
```

```
-----
/scratch usage for pawsey0001, used = 17.85 TiB, limit = 2048.00 TiB
/software usage for pawsey0001, used = 166.66 GiB, limit = 256.00 GiB
```

OUTPUTS: Monitoring Filesystem Usage

```
$ lfs quota -h -g $PAWSEY_PROJECT /scratch
```

```
Disk quotas for grp pawsey0001 (gid 30110):
```

Filesystem	used	quota	limit	grace	files	quota	limit	grace
/scratch	17.85T	0k	2P	-	588007	0	0	-

```
$ lfs quota -h -u $USER /scratch
```

```
Disk quotas for usr username (uid userid):
```

Filesystem	used	quota	limit	grace	files	quota	limit	grace
/scratch	590.5M	0k	0k	-	3393	0	2000000	-

```
$ lfs quota -h -g $PAWSEY_PROJECT /software
```

```
Disk quotas for grp pawsey0001 (gid 30110):
```

Filesystem	used	quota	limit	grace	files	quota	limit	grace
/software/	166.7G	0k	256G	-	700083	0	0	-

```
$ lfs quota -h -u $USER /software
```

```
Disk quotas for usr username (uid userid):
```

Filesystem	used	quota	limit	grace	files	quota	limit	grace
/software	6.84G	0k	0k	-	12987	0	100000	-

Summary



- Terms we learnt
 - Filesystem



- Tasks we learnt
 - Monitor filesystem quotas: `quota` and `pawseyAccountBalance`



- Avoid keeping workflow-related settings inside shell configuration files



- Be mindful of the 30-day purge policy on `/scratch`
- Proactively monitor filesystem usage, both in size and file count



Getting Help



Australian Government



Getting Help

<https://support.pawsey.org.au>

Pawsey has extensive [User Support Documentation](#).

Areas covered include:

- System user guides
- Knowledge Base
- Pawsey-supported software list
- Maintenance logs
- Policies and terms of use

For further assistance, contact the help desk, via [User Support Portal](#).

Help us to help you by providing details, such as:

- Which resource
- Error messages
- Location of files
- SLURM job id
- Your username if having login issues
- Never tell us (or anyone) your password!

Become a Pawsey Friend and receive our Newsletter:

<https://pawsey.org.au/pawsey-friends/>



Q & A Session



Australian Government

