

ZWS at Pawsey

Advancing Science through Supercomputing


The Pawsey Supercomputing Centre is an unincorporated joint venture between




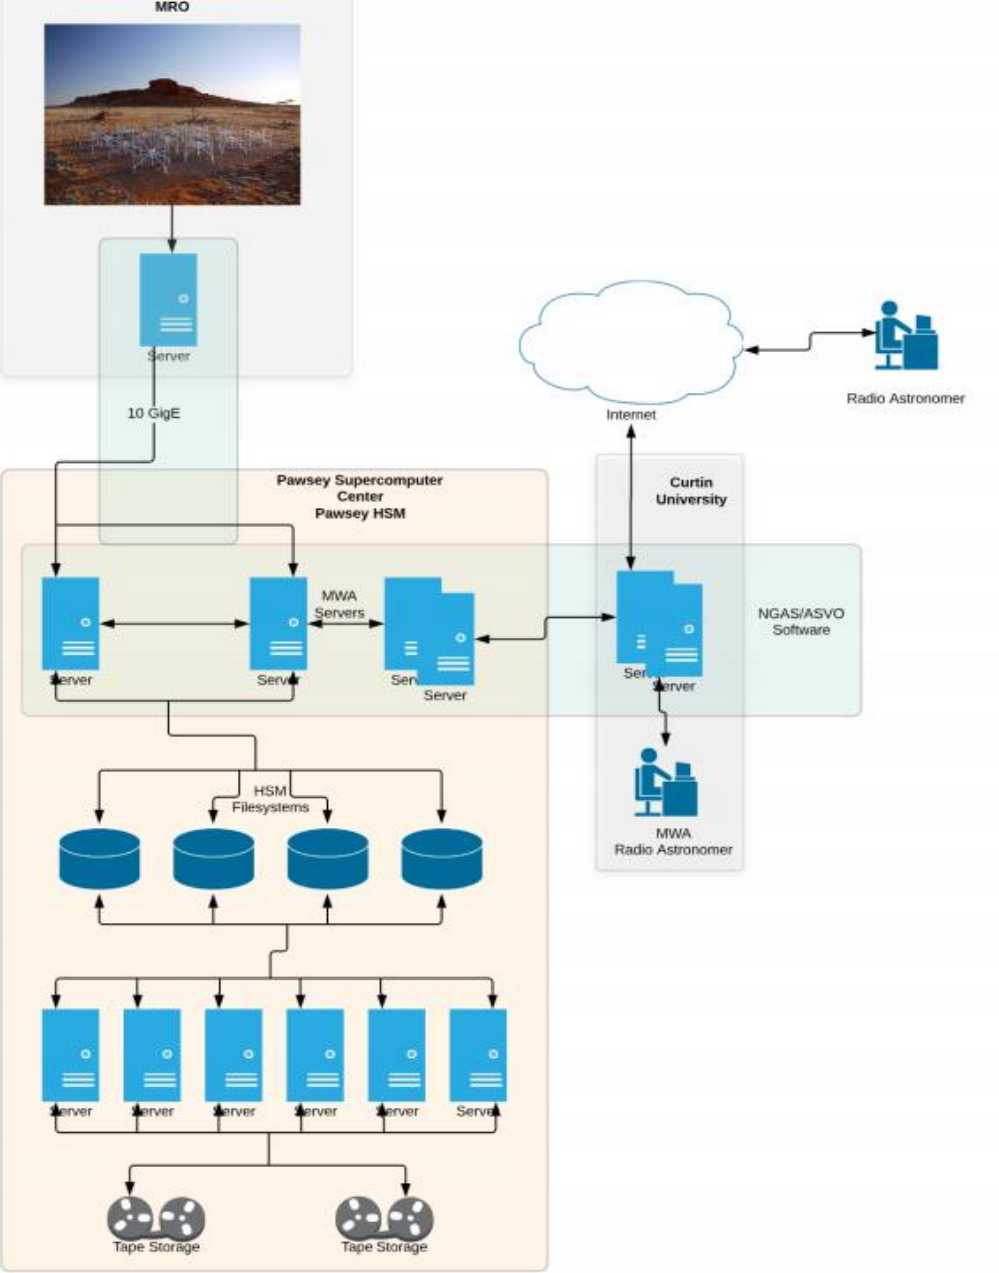
and proudly funded by




Pawsey Supercomputing Centre




MWA






Observations

- Data stream split into 24 or 48 files
- 1 GB or so in size




File sizes, mw01fs

Digits	Size from	to < than	Num Files
0	0	0	564
1	1.0B	10B	11
2	10B	100B	1810
3	100B	1.0kB	5655
4	1.0kB	10kB	9539
5	10kB	100kB	239318
6	100kB	1.0MB	50387
7	1.0MB	10MB	65126
8	10MB	100MB	97499
9	100MB	1.0GB	2204437
10	1.0GB	10GB	3781097
11	10GB	100GB	4464
12	100GB	1.0TB	16




Data Access

- Accessing one file in an observation recalls all 24 or 48 files
- On average, 3 mins to load a cartridge
- Each 1GB file read in 3-5 seconds
- Processing of an observation can't start until the last file is recalled
- Tape drive, cartridge, system load, ..., cause delays
- Faster tape drives will make things worse




The ZWS Solution

- 3.6PB of ZWS
- 6 x D6020
- 2 x pDMOs
- Configured as Fast Mount Cache
- vg1_j00 and vg2_j00, round robin
- Only for MWA at present




Does it Work?

- Testing, repeatedly recalled 48 files from;
- Tape: 04:07 – 316:54; busy system, no free drives
- ZWS: 0:42 – 1:23; pre-production, sole use



In Practice?

- 33PB of MWA data
- In 5 filesystems, 1.7PB total size
- 3.6PB of ZWS
- Recalls;
 - Tape: 10TB, 14173 files
 - 7TB, 4001 files
- ZWS: 10.8TB, 42412 files
 - 10.1TB, 39202 files
- 50% or so hit rate. It does vary



The future?

- More ZWS or similar
- Recall prediction
- Schedule recalls in advance

Find out more

Pawsey Website (www.pawsey.org.au)

Pawsey Friends mailing list

Pawsey Twitter feed (@PawseyCentre)

User Support Portal (support.pawsey.org.au)



