



# CSIRO Scientific Computing Site Report

Peter Edwards | Sr Systems & Data Administrator  
DMF Users Group | August 2019

CSIRO IM&T SCIENTIFIC COMPUTING  
[www.csiro.au](http://www.csiro.au)



# What is the Commonwealth Scientific and Industrial Research Organisation?

The *Science and Industry Research Act 1949* defines our purpose and the functions we undertake for the benefit of Australia:

- To carry out scientific research for any of the following purposes:
  - Assisting Australian industry;
  - Furthering the interests of the Australian community;
  - Contributing to the achievement of Australian national objectives or the performance of the national and international responsibilities of the Commonwealth; and
  - Any other purpose determined by the Minister;
- To encourage or facilitate the application or utilisation of the results of such research.

# What is CSIRO doing?

With more than 1,800 patents, we are Australia's largest patent holder. This ever-increasing wealth of intellectual property is a vast source of commercial opportunity and has already resulted in more than 150 spin-off companies.

Even just considering six CSIRO contributions, the Australian economy reaps \$5 billion a year in benefits from

- the Australian Animal Health Laboratory
- our work in cotton
- our longwall mining technology
- our Opticool energy control system
- our Novacq prawn feed
- our water resource assessment work.

# What has CSIRO done?



# How big is CSIRO?

We have more than 5,000 experts based in 55 centres and extensive local and international networks.

We collaborate with 3,000 customers each year, including

- Australian federal, state and local government bodies
- small, medium and large businesses
- the majority of Research Development Corporations, Cooperative Research Centres and Australian universities
- more than 150 international partners.

# RV Investigator



# Not to be confused with:



# CSIRO BeeGFS Scratch Filesystem

- Tendered for 2PB – Dell EMC/Pacific Tech won the bid
- Replace aging NFS solution designed for 128 node cluster
  - Now ~500 nodes
- Tweaked configuration for different NVMe IO characteristics
- NVMe drive supply was a problem
- Who is Pacific Tech?
  - Gold Partner of ThinkParq/BeeGFS and Systems Integrator
  - On the Exhibition Floor – go and check them out
- Who is Dell EMC? 😊 – also on the Exhibition Floor.





# Hardware

- 1 management server
  - 6 x 600GB SAS
- 4 metadata servers
  - 24 x 1.6TB Intel P4600 NVMe
    - 560K/177K R/W IOPS
  - Dual 3.0GHz 12core, 384GB
  - Dual ConnectX-5 EDR
- 32 storage servers
  - 24 x 3.2TB Intel P4600 NVMe
    - 695K/226K R/W IOPS
  - Dual 2.4GHz 20core, 192GB
  - Dual ConnectX-5 EDR



# Status

- Bright Image built and refreshed with BeeGFS 7.1
- Bright configuration work underway
- Pacific Tech primed for pass 1 benchmarking
  - mdadm vs ZFS for storage servers
  - SW RAID configuration choices
    - 2 x 10+2 or 22+2 etc.
  - IO-500 run
    - **Came in at number 6!**
- Go production with half of it for scratch
- Pass 2 benchmarking
- Maybe provide a portion for Machine Learning training data
- 2<sup>nd</sup> management server for HA

# DMF stats (May 2018)

	units	ruby(*)	dmfact03	dmfact04	dmfact05	dmfvic02	dmfvic03
Purpose		data	rsync	data	data	data	rsync
Data managed	TB	11,625	4,835	1,438	1,729	248	2,432
Inodes	M	198	2,435	160	0.834	9.1	2,120
Mig files	M	60	406	154	0.291	8.1	268
Data mig	TB/mth	373	459	87	66	104	39
Files mig	k/mth	2,432	19,328	346	107	7.5	1,783
Files del/mod	k/mth	1,440	12,452	0	0.001	110	5,109
Data rcl	TB/mth	860	2.3	37	43	1.7	2.2
Files rcl	k/mth	524	11	15	0.426	0.059	19
Tape mnts	/mth	14,921	7,010	869	327	183	4,857

(\*) Just the user accessible data filesystem shown, 4 rsync filesystems omitted

# Thank you

**CSIRO IM&T Scientific Computing Services**  
Peter Edwards

**t** +61 3 9545 2377

**e** [peter.edwards@csiro.au](mailto:peter.edwards@csiro.au)

**w** [www.hpsc.csiro.au](http://www.hpsc.csiro.au)

**CSIRO IM&T SCIENTIFIC COMPUTING SERVICES**  
[www.csiro.au](http://www.csiro.au)



# Computing at CSIRO

- 1949 - CSIR Mk1 - the fifth computer in the world, weighing 7 tonnes, using 30kW of power for its 2000 valves (vacuum tubes), initially with 768 20-bit words of mercury acoustic delay line RAM. Later renamed to CSIRAC.
  - <https://csiropedia.csiro.au/the-computer-csirac-1965/>
  - <https://en.wikipedia.org/wiki/CSIRAC>
- 1955 - CSIRO decides to pursue cloud seeding rather than computing, so CSIRAC transferred to University of Melbourne.
- Time passes...
- 1990 - Cray Research Y-MP2/216 (Cherax #1)
- 1991 – DMF 2.0 or 2.1 with manually mounted 3480 tape cartridges
- 1993 - StorageTek 4400 tape library

# Cherax #1

- Cray Y-MP2/216 – “big” (1.2 GB), expensive, fast disk
  - Turned a compute problem into a storage problem!



# HSM at CSIRO

- 1 combined HPC system & DMF server (UV3000) in Canberra with:
  - hybrid filesystems with quotas
  - Copan MAID
  - DCM
  - STK T10kC/D tapes, plus some T10kB still being phased out
  - FCoIP link to interstate tape drives
- Same DMF-managed \$HOME filesystem since 1991
- Compute and DMF functions are to be separated from each other
- 3 dedicated DMF servers in Canberra (to become 2)
- 2 dedicated DMF servers in Melbourne (to become 1)
- 10 other HSM servers
  - 2 each in Canberra, Melbourne, Brisbane, Perth & Hobart (to become 1 each)