ANITA Proposal for a National astronomy compUTing Strategy (PNUTS)

AAL data and computing meeting 23rd Oct 2018

Daniel Price, on behalf of ANITA committee and Australian theoretical astrophysics community (ANITA membership)

tiny.cc/anitareport

Decadal Plan

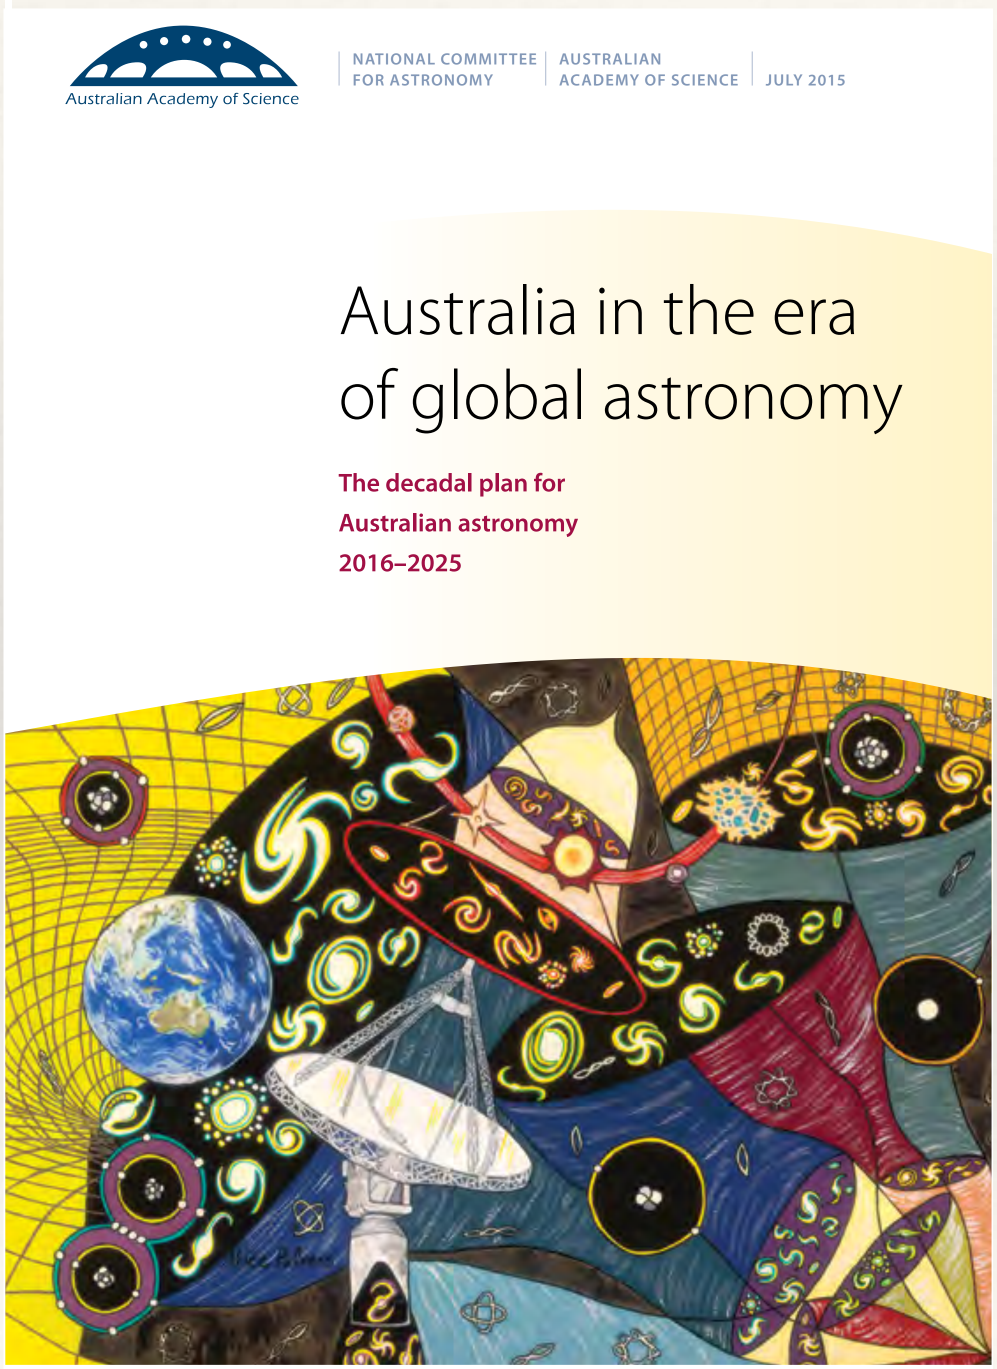

1 Executive summary

“ASTRONOMY IS ENTERING A NEW ERA OF DISCOVERY WITH GLOBAL FACILITIES”

From the birth of the first stars and galaxies to the origin of habitable planets, Australian astronomy will continue to make world-leading discoveries over the coming decade. These discoveries will be achieved through a strong platform of cutting-edge optical/infrared and radio telescopes, supported by theoretical and computational astrophysics. This Decadal Plan identifies five top-level science infrastructure priorities. These priorities are equally weighted as part of an overall astronomy capability:

- Partnership equating to 30% of an 8-metre class optical/infrared telescope;
- Continued development and operations of Square Kilometre Array (SKA) precursors, the Australian SKA Pathfinder (ASKAP) and Murchison Widefield Array (MWA) at the Murchison Radio-astronomy Observatory (MRO), and membership of the SKA telescope;
- Partnership equating to 10% of a 30-metre class optical/infrared extremely large telescope (ELT), such as the Giant Magellan Telescope (GMT);
- Capability within the national observatories (the Australian Astronomical Observatory, AAO; and Australia Telescope National Facility, ATNF) to maximise Australia's engagement in global projects through instrumentation development for these and other facilities;
- World-class high performance computing (HPC) and software capability for large theoretical simulations, and resources to enable processing and delivery of large data sets from these facilities.

1 EXECUTIVE SUMMA

Figure 4.2: Graphical summary of the citation-weighted impact of different techniques in Australian astronomy in the decade to 2014 (bottom), showing the relative changes since 2005 (top). Optical/IR refers to use of the combination of ground-based optical/IR, and space-based optical/IR and UV facilities

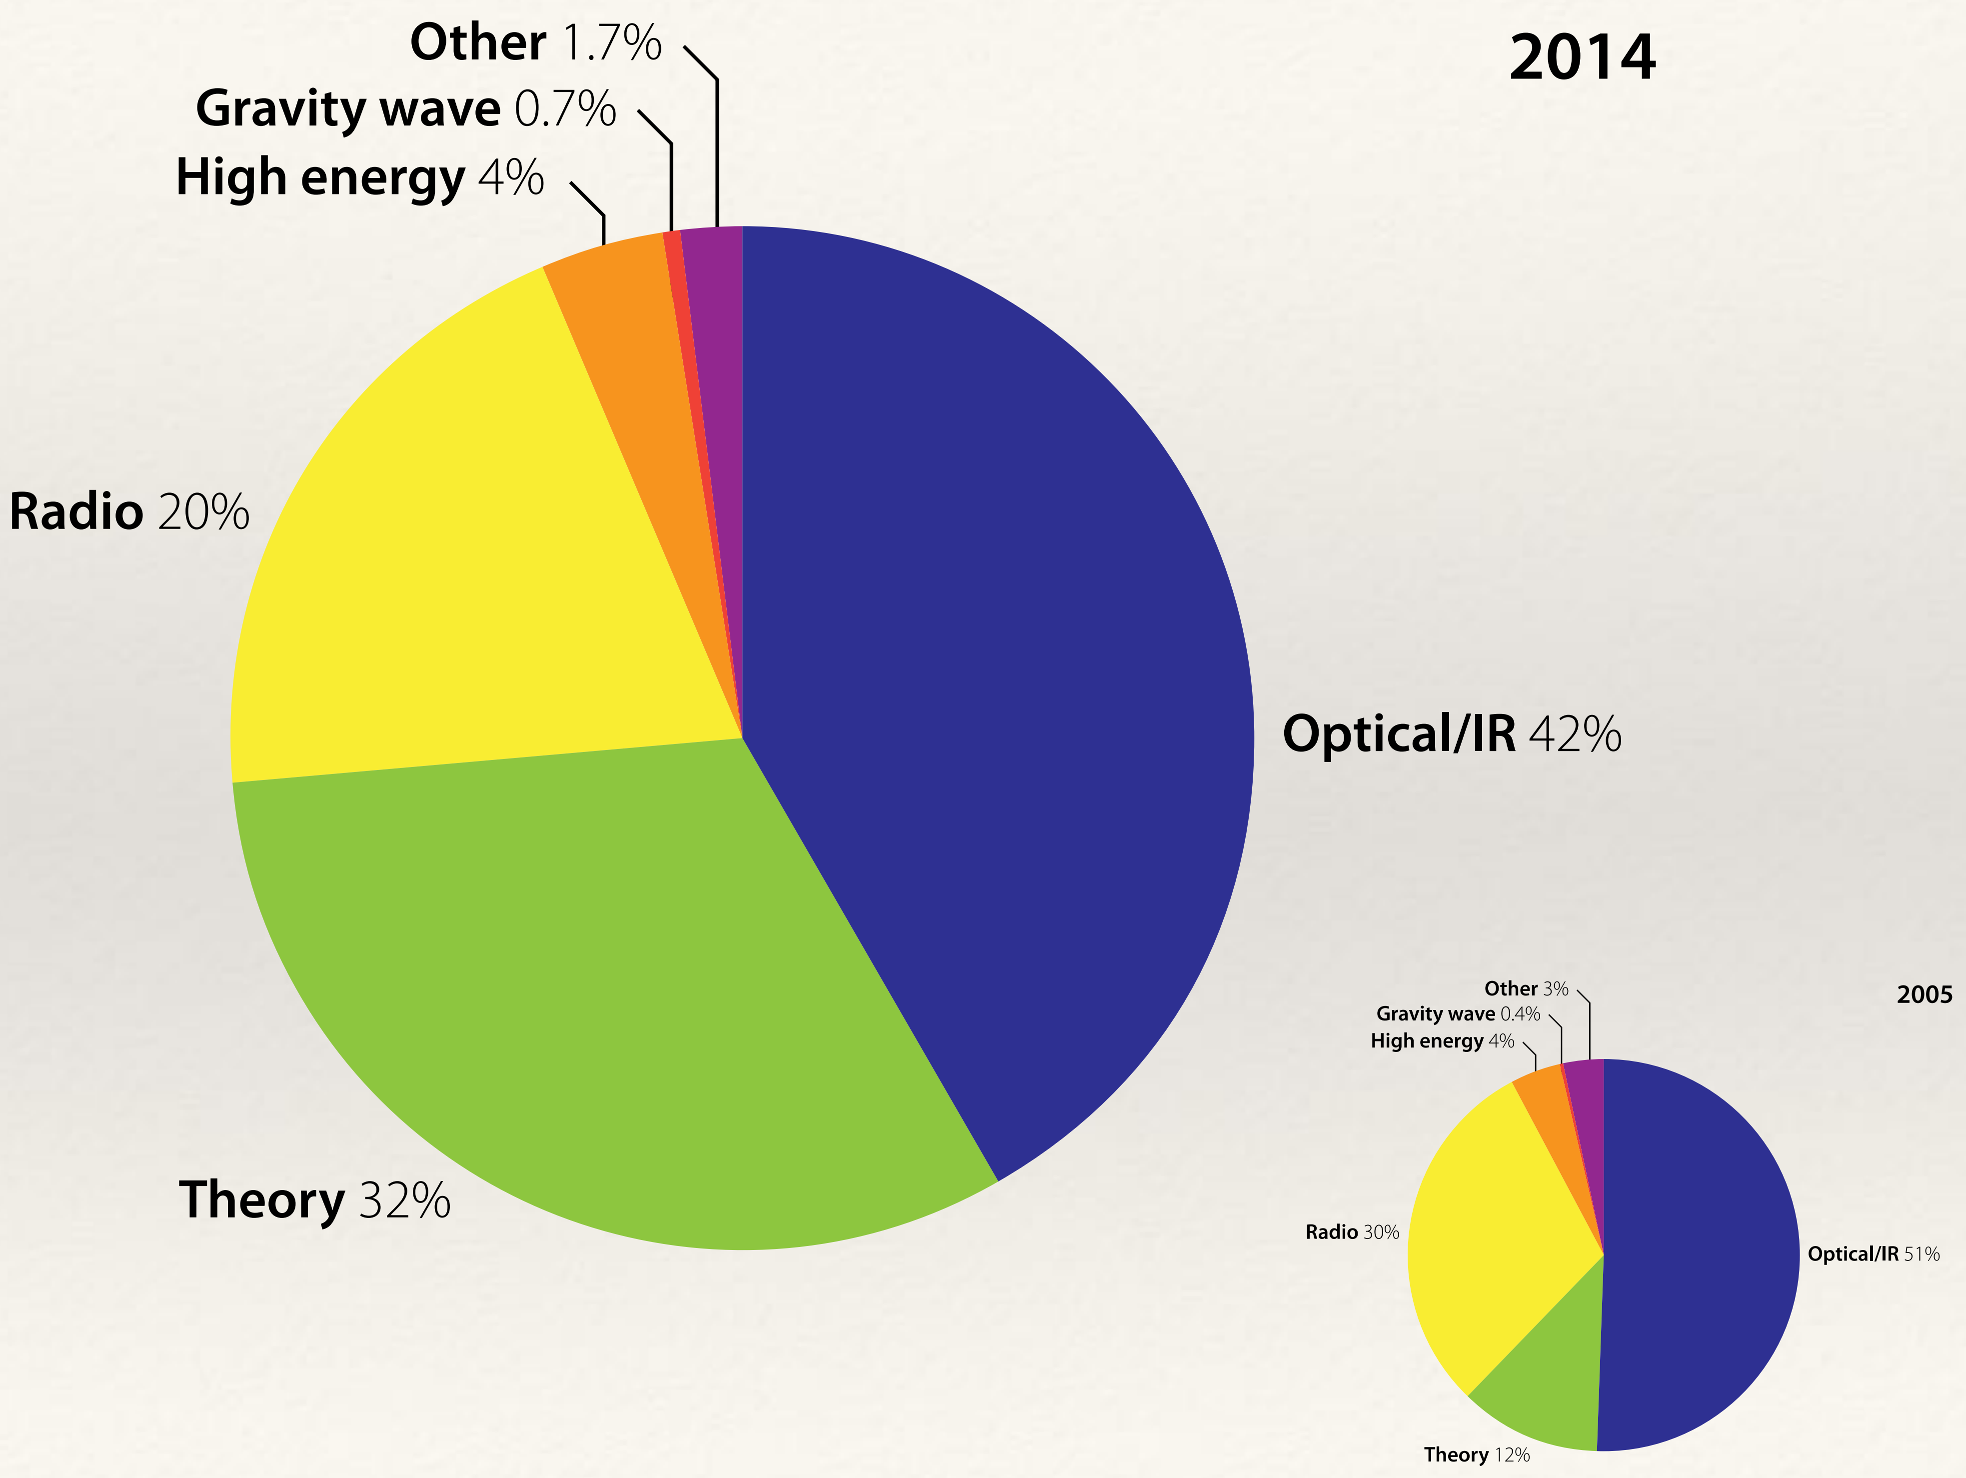

Figure 4.3: Techniques used by Australian astronomers, expressed as the percentage of researcher effort. Many individual astronomers contribute across multiple techniques. The wedge labelled 'Other' includes astronomical areas not elsewhere listed

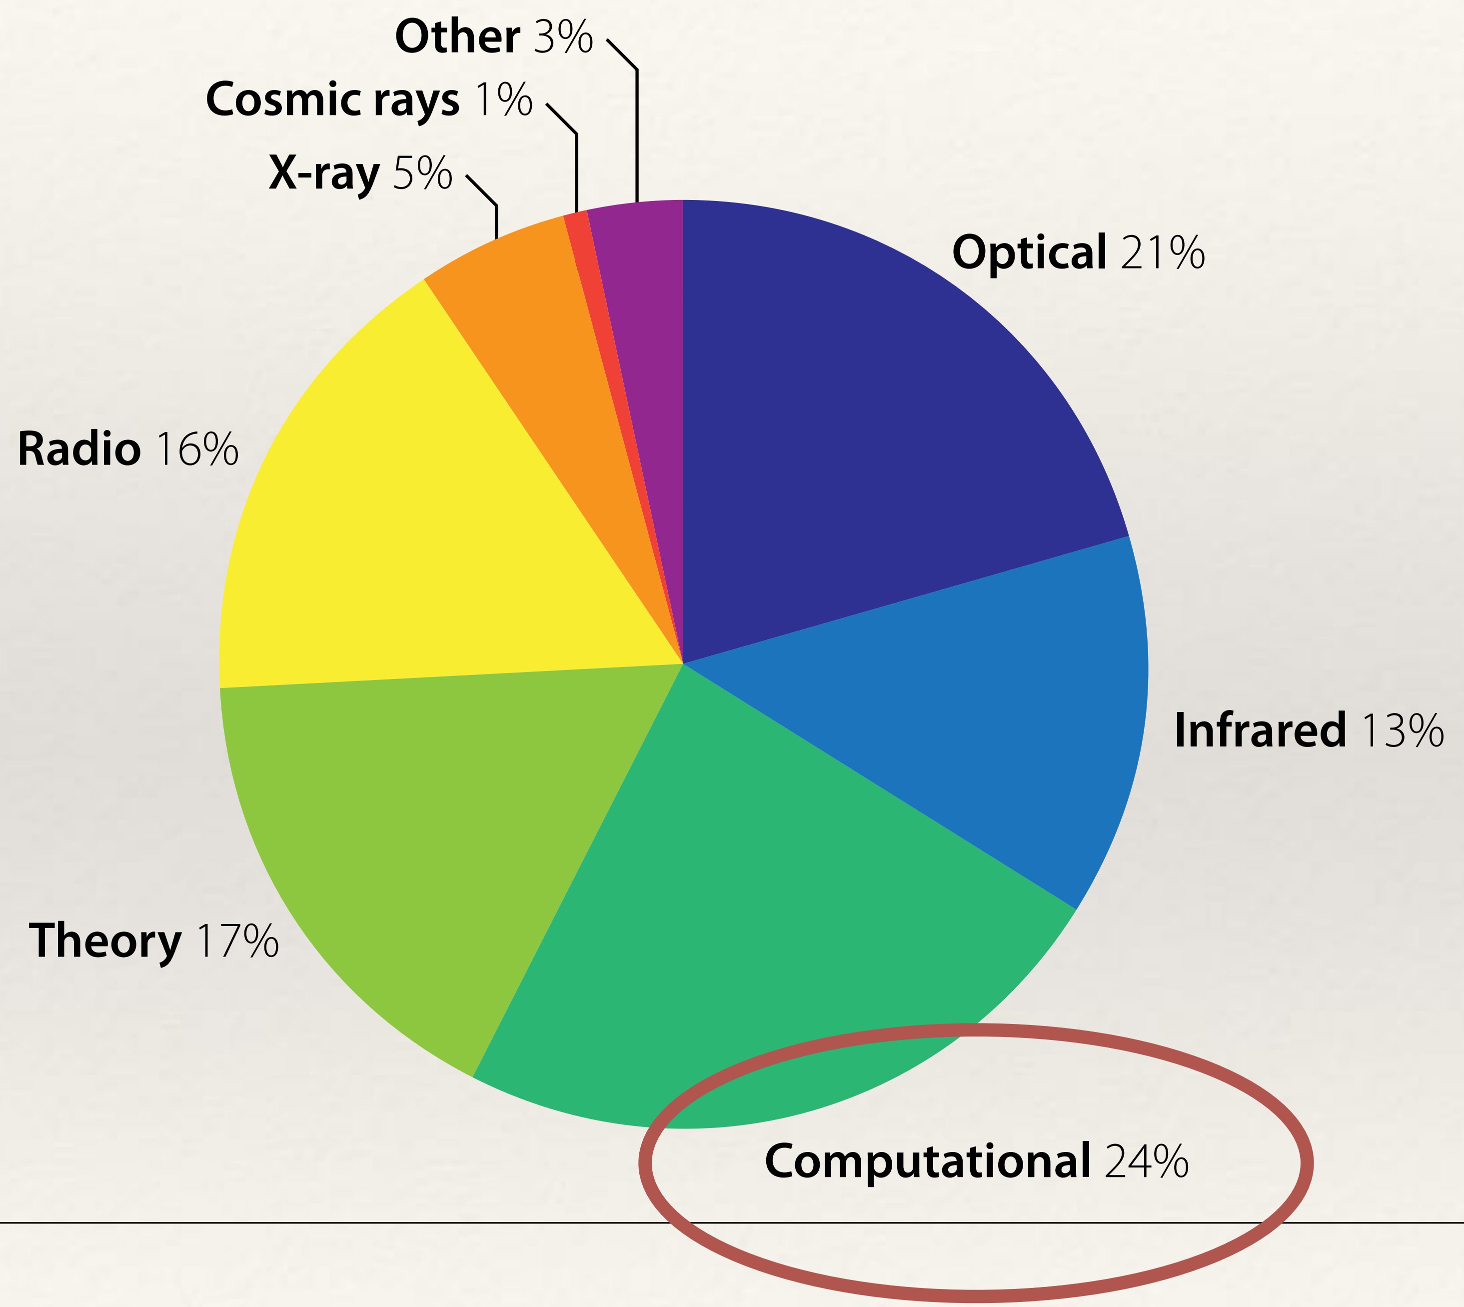

Decadal plan priorities (p1)

- Partnership equating to 30% of an 8-metre class optical/infrared telescope;
- Continued development and operations of Square Kilometre Array (SKA) precursors, the Australian SKA Pathfinder (ASKAP) and Murchison Widefield Array (MWA) at the Murchison Radio-astronomy Observatory (MRO), and membership of the SKA telescope;
- Partnership equating to 10% of a 30-metre class optical/infrared extremely large telescope (ELT), such as the Giant Magellan Telescope (GMT);
- Capability within the national observatories (the Australian Astronomical Observatory, AAO; and Australia Telescope National Facility, ATNF) to maximise Australia's engagement in global projects through instrumentation development for these and other facilities;
- World-class high performance computing (HPC) ~~and software capability~~ for large theoretical simulations, and ~~resources to enable processing and delivery of large data sets~~ from these facilities.

Decadal plan - p5

Theoretical and computational astrophysics has grown to become a focus across all areas of strength in Australian astronomy research, now representing approximately one-third of its research impact. Alongside the growing need for HPC to process data products from the next generation of telescopes, **computational astrophysics in Australia will also require strategic HPC investment**.

"The estimated HPC resource needed for Australian astronomy to achieve these goals corresponds to approximately the equivalent of 30% of a top-100 supercomputer."

AAL funding of data and computing so far

- Data and software: Funding for ASVO nodes + ADACS, both good initiatives
- Computing: OzSTAR received $180K contribution from AAL - **factor of 5 drop** from $1m via EIF towards original gSTAR
- ANITA expressed concern to AAL (29th March letter) that only data and software strategy so far
- Need strategy for computing investment

Requirements for a computing strategy

- **Flexible**. Not "one size fits all"
- **High throughput**
- **Internationally competitive**
- **Low barrier to entry**

Flexibility: Proposed Tier Structure

Tier level	Description	Typical workflow	Interconnect	Example facilities	Astronomy applications	Access mechanism
Tier 0	National centres open to all disciplines	Large individual jobs running on thousands of CPUs	Tightly coupled/ infiniband or similar	NCI, Pawsey	Cosmological simulations, supernovae or star formation simulations using MPI parallel fluids codes	Merit allocation scheme, ideally twice per year
Tier 1	National level but specialised facilities (e.g. specific to Astronomy)	Parameter studies using tens to hundreds of CPUs per job	Tightly coupled/ infiniband or similar	gSTAR, OzSTAR, Astronomy partner share in national facilities	Fluid simulations using MPI/OpenMP parallel codes. N-body calculations using GPUs.	Available on demand, with large allocations by merit allocation
Tier 2	Grid / cloud computing	Large number of single CPU jobs	No communication between jobs	Nectar grid	Parameter estimation, Monte-Carlo Markov chain, Bayesian inference, Stellar evolution codes	Available on demand

tiny.cc/anitareport

High throughput computing

- Work achieved over month-long timescales, not just time for one job
- Jobs start running within 24 hours
- Minimise downtime for maintenance / disk failures
- Availability of machines appropriate to the workflow. Requires dedicated Astronomy-focussed Tier 1 and Tier 2 facilities
- Low startup cost - no need for merit allocation on lower tiers, can just "join the queue". Fairness established by job scheduler.
- More frequent merit allocation (semi-annual or rolling basis)

"Raijin just couldn't cope - too many files, too many single core jobs - after trying to find a solution for several months … was so inefficient that the stuff I got the time to do never got done. A bit of a tale of woe, really - but the key point is the flexibility to work with users with unusual use cases is key - and to be able to cope with moderately large numbers of files and jobs, of course"

tiny.cc/anitareport

Internationally competitive?

- Australia has only **one** machine in top 100 supercomputers (Raijin at NCI, followed by Magnus at Pawsey #217)
- 115M cpu hours available (in total) as open time to *entire* scientific community, ~11% of this to Astronomy in last NCMAS call

"The [NCMAS] Committee will consider allocation of more than 4-5 MSU/year only to applications which demonstrate exceptional …"

tiny.cc/anitareport

Internationally competitive?

System	Architecture	Site (Country)	Core Hours *(node hours)*	Minimum request
Curie	Bull Bullx cluster	GENCI@CEA (FR)	**128 million** *(5.5 million)*	**15 million core hours**
Curie successor	BULL Sequana (starting service in Q2 2018)	GENCI@CEA (FR)	**36 million** *(0.5 million)*	**15 million core hours**
Hazel Hen	Cray XC40 System	GCS@HLRS (DE)	**57 million** *(2.4 million)*	**20 million core hours**
Marconi-Broadwell	Lenovo System	CINECA (IT)	**36 million** *(1 million)*	**15 million core hours**
Marconi-KNL	Lenovo System	CINECA (IT)	**442 million** *(6.5 million)*	**30 million core hours**
MareNostrum	Lenovo System	BSC (ES)	**475 million** *(10 million)*	**15 million core hours**
Piz Daint	Cray XC50 System	CSCS (CH)	**510 million** *(7.5 million)*	**68 million core hours Use of GPUs**
SuperMUC	IBM System X iDataplex/ Lenovo NextScale	GCS@LRZ (DE)	**44 million** *(2.4 million)*	**15 million core hours**

15th Call for Proposals for Project Access

The PRACE 15th Call for Proposals for Project Access was open until 30 May 2017

"The real big problem in Australia is the poor Tier 0 facilities, which essentially preclude participation in certain computational fields (SNe modelling, large-scale cosmological simulations) by Australians who don't have some way of getting access to time overseas. I did a survey last year, before the AAL AeRAC ceased to exist, and found that >50% of the CPU time being used by Australians was outside Australia, and that is in itself a biased sample because we just don't have certain fields represented in the country due to the lack of computer facilities."

Missing pieces

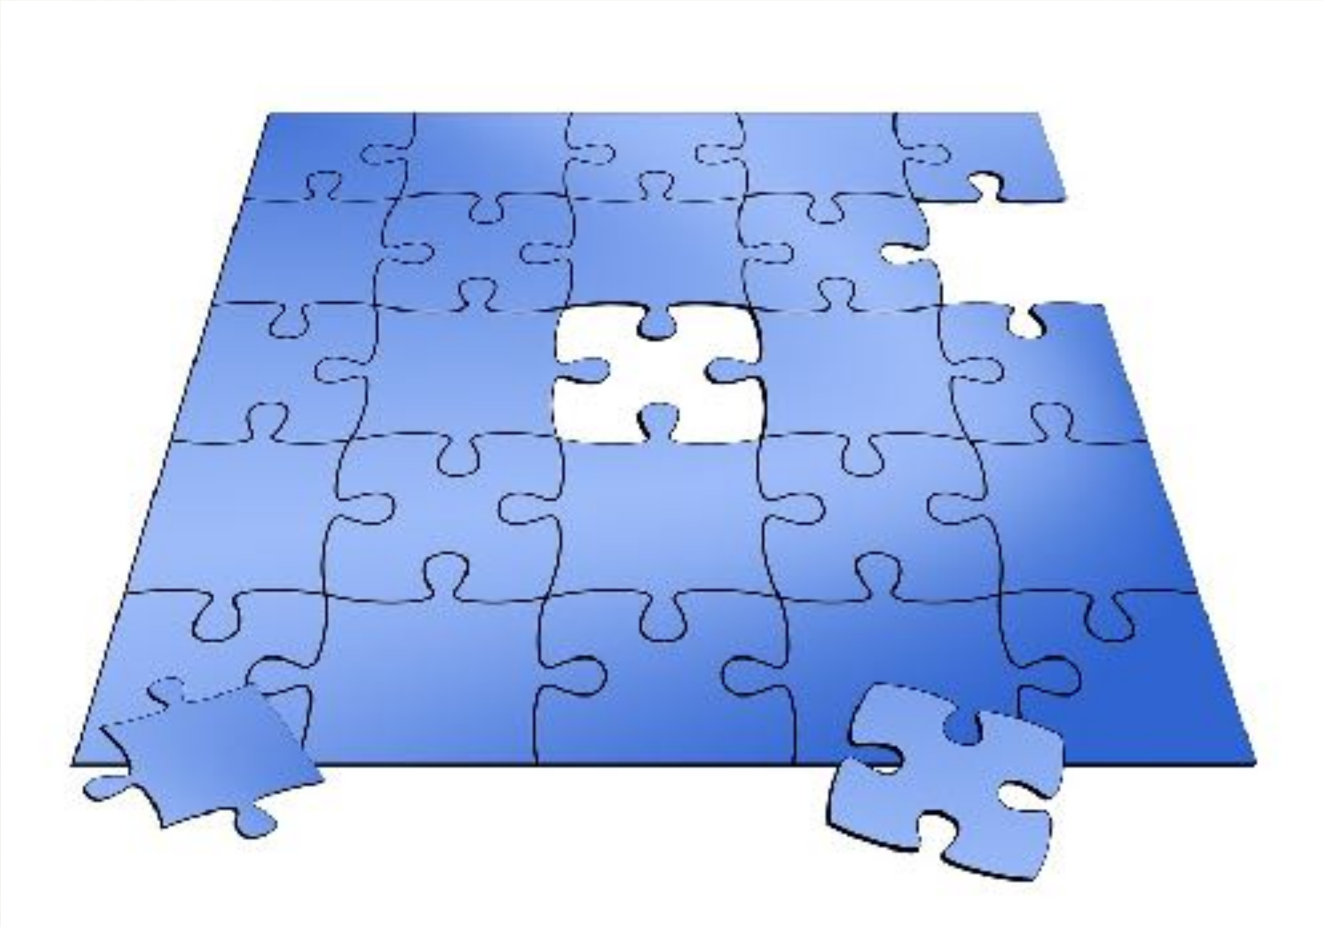

- No internationally competitive Tier 0 facility
- No easily accessible Tier 2 grid computing facility

ANITA recommendations

- AAL to fund, co-fund, or secure use of, on average, continuous use of at least 5000 VCPUs as a Tier 2 cloud computing facility for Australian Astronomy
- Plan to upgrade OzSTAR on 3-5 year timescale (Tier 1)
- Secure access to internationally competitive Tier 0 facilities, with minimum 200M cpu hours (2017 equivalent) available for Australia Astronomy annually. Could we join EU PRACE?
- Semi-annual merit allocation calls
- Prioritise specialised software support via ADACS

tiny.cc/anitareport