


Q1 Could you please indicate your research area(s) (you could choose more than one research area)?

Answered: 65 Skipped: 0




| ANSWER CHOICES | RESPONSES |
|------------------------|-----------|
| Computational | 50.77% 33 |
| Gravitational Wave | 12.31% 8 |
| High energy | 7.69% 5 |
| UV/Optical/IR | 23.08% 15 |
| Radio | 36.92% 24 |
| Theory | 43.08% 28 |
| Other (please specify) | 16.92% 11 |
| Total Respondents: 65 | |

| # | OTHER (PLEASE SPECIFY) | DATE |
|----|--|--------------------|
| 1 | Instrumentation | 4/17/2019 9:17 AM |
| 2 | Magnetohydrodynamical simulations of the interstellar medium | 4/16/2019 5:36 PM |
| 3 | Exoplanets; Solar System | 4/16/2019 3:49 PM |
| 4 | protoplanets and protoplanetary discs | 4/16/2019 3:04 PM |
| 5 | Galaxy Evolution | 4/16/2019 2:55 PM |
| 6 | Star Formation | 4/16/2019 2:43 PM |
| 7 | Astrophysical stellar transients | 4/16/2019 2:36 PM |
| 8 | galaxy formation | 4/16/2019 1:43 PM |
| 9 | Cosmology | 4/16/2019 1:04 PM |
| 10 | Cosmology, cmb | 4/16/2019 12:50 PM |
| 11 | machine learning + natural language | 4/16/2019 12:39 PM |

Q2 This question refers to the HPC outputs of others (either your collaborators or the community). How heavily do you rely on the HPC outputs of others for your research?


Answered: 65 Skipped: 0



| ANSWER CHOICES | RESPONSES |
|---------------------|-----------|
| Critical | 43.08% 28 |
| Important | 36.92% 24 |
| Occasionally Useful | 16.92% 11 |
| Not Useful | 3.08% 2 |
| TOTAL | 65 |

Q3 This question refers to your own personal HPC usage. Could you please indicate your current HPC usage (e.g. CPU/GPU hours)?


Answered: 65 Skipped: 0



| ANSWER CHOICES | RESPONSES | |
|---------------------|-----------|-----------|
| <1M hours/year | 44.62% | 29 |
| 1M - 5M hours/year | 38.46% | 25 |
| 5M - 10M hours/year | 9.23% | 6 |
| > 10M hours/year | 7.69% | 5 |
| TOTAL | | 65 |

Q4 What sort of HPC architecture suits your research best? If a mixture, please estimate how much of each.

Answered: 65 Skipped: 0




| ANSWER CHOICES | RESPONSES |
|-----------------------|-----------|
| CPU | 89.23% 58 |
| GPU | 24.62% 16 |
| Cloud | 12.31% 8 |
| Fast input-output | 24.62% 16 |
| Memory intensive | 47.69% 31 |
| Other | 3.08% 2 |
| Total Respondents: 65 | |

| # | IF A MIXTURE, PLEASE ESTIMATE HOW MUCH OF EACH. | DATE |
|----|--|-------------------|
| 1 | 30% CPU, 50% GPU, 20% Cloud | 4/25/2019 6:32 PM |
| 2 | 80:20 | 4/25/2019 4:11 AM |
| 3 | We regularly use e.g. 240 CPUs with 500 GB RAM and a fast I/O to ensure a high data rate of processing | 4/18/2019 1:00 PM |
| 4 | Majority on AWS. Most radio is IO limited | 4/17/2019 7:15 PM |
| 5 | Roughly 50:50, with about 80% being memory intensive | 4/17/2019 9:17 AM |
| 6 | Currently probably 90/10 CPU/GPU, but GPUs likely to become increasingly important | 4/17/2019 6:27 AM |
| 7 | Fluctuates. Currently mostly multiple CPUs on one node. | 4/17/2019 1:58 AM |
| 8 | usually > 1024 CPU cores | 4/16/2019 5:36 PM |
| 9 | High-throughput computing, i.e. large numbers of independent jobs | 4/16/2019 3:47 PM |
| 10 | Hard to say: for exoplanets hydro-sims we equally require large numbers of CPU cycles and large quantities of RAM (so say 50% + 50%) | 4/16/2019 3:04 PM |

| | | |
|----|--|--------------------|
| 11 | 50Gb RAM per CPU; 50 CPUs for several days; 1-2 runs per month | 4/16/2019 2:00 PM |
| 12 | Memory intensive is critical for the analysis of the simulations, while CPU is necessary to produce the simulations | 4/16/2019 1:43 PM |
| 13 | 80% cpu, 10% fast i/o, 10% memory intensive | 4/16/2019 1:19 PM |
| 14 | Experimenting with GPU, but that's not in production mode yet, so almost all usage is CPU. | 4/16/2019 1:11 PM |
| 15 | 70/30 CPU/Memory Intensive. Most jobs can be run with reasonable amounts of memory (~4GB/cpu), however there are often times when much more is critical. | 4/16/2019 1:04 PM |
| 16 | Data processing accounts for 50-70% of my work and depending on the pipelines this is a mixture of CPU and GPU. But certainly all of it is memory intensive. For other side projects, cloud and GPU computing is important at the 20-30% level | 4/16/2019 12:57 PM |
| 17 | CPU 60% Memory intensive 40% | 4/16/2019 12:56 PM |
| 18 | 50-50 | 4/16/2019 12:55 PM |
| 19 | Note by cloud, I meant grid computing. | 4/16/2019 12:50 PM |
| 20 | 80/20 CPU/GPU | 4/16/2019 12:41 PM |

Q5 This question refers to your own personal HPC usage. Where do your current HPC resources come from? If a mixture, please estimate how much of each.

Answered: 65 Skipped: 0




| ANSWER CHOICES | RESPONSES | |
|---------------------------------|-----------|----|
| NCI | 38.46% | 25 |
| Pawsey | 29.23% | 19 |
| OzSTAR | 35.38% | 23 |
| Other Domestic (please specify) | 26.15% | 17 |
| Other Overseas (please specify) | 24.62% | 16 |
| Total Respondents: 65 | | |

| # | PLEASE SPECIFY, IF YOU CHOOSE "OTHER DOMESTIC" OR "OTHER OVERSEAS" | DATE |
|----|---|-------------------|
| 1 | University facility | 4/29/2019 3:54 PM |
| 2 | local custers; cluster at MSU | 4/25/2019 4:11 AM |
| 3 | local university HPC facilities | 4/19/2019 9:40 AM |
| 4 | Monash | 4/18/2019 2:33 PM |
| 5 | Our own 64-CPU server | 4/18/2019 1:00 PM |
| 6 | AWS (USA) | 4/17/2019 7:15 PM |
| 7 | NERSC (US), U.Chicago (US), Local University Cluster | 4/17/2019 8:48 AM |
| 8 | 20% Raijin, 80% Ozstar | 4/17/2019 1:58 AM |
| 9 | SuperMUC for accessing petabyte scale simulations. | 4/16/2019 5:36 PM |
| 10 | EngineRoom, Azure, AWS | 4/16/2019 4:53 PM |
| 11 | the university | 4/16/2019 3:56 PM |
| 12 | USQ's Fawkes cluster | 4/16/2019 3:49 PM |
| 13 | LSC DataGrid: https://www.lsc-group.phys.uwm.edu/lscdatagrid/ | 4/16/2019 3:47 PM |
| 14 | DIRAC (UK) and Rocket (Newcastle University) | 4/16/2019 3:43 PM |

| | | |
|----|---|--------------------|
| 15 | CSIRO | 4/16/2019 3:34 PM |
| 16 | Monash sun grid | 4/16/2019 3:03 PM |
| 17 | 40% NCI, 40% AVATAR at RSAA, ANU, 20 UHHPC at University of Hertfordshire UK | 4/16/2019 2:55 PM |
| 18 | NCI: 10%, Kunanyi (University of Tasmania): 90% | 4/16/2019 2:45 PM |
| 19 | Leibniz Supercomputing Centre | 4/16/2019 2:43 PM |
| 20 | 90% NCI, 10% in-house cluster | 4/16/2019 1:52 PM |
| 21 | I use computers in the UK as well for analysis as accessing nodes with large memory in Australia is difficult (with >300GB of memory) | 4/16/2019 1:43 PM |
| 22 | XSEDE | 4/16/2019 1:19 PM |
| 23 | NASA Ames -- Pleiades | 4/16/2019 1:11 PM |
| 24 | ICRAR/UWA | 4/16/2019 1:04 PM |
| 25 | Other overseas: UK-based supercomputers (Sciama, Cosma). Mainly storage. Over 90% of usage this year has been on OzStar. In previous years this mostly came from NCI and Pawsey, but oversubscription (and lack of memory on Pawsey) has led to an increased usage of OzStar. | 4/16/2019 1:04 PM |
| 26 | Most resources accessed through joint applications (with European colleagues) to PRACE or DIRAC. Equivalent resources are difficult to obtain in Australia. | 4/16/2019 12:59 PM |
| 27 | We use a bit of local computing at UoM. Mostly NERSC in the US, and Open Science Grid. | 4/16/2019 12:50 PM |
| 28 | 30/70 Ozstar/LIGO data grid | 4/16/2019 12:41 PM |
| 29 | Local University machines | 4/16/2019 12:37 PM |

Q6 This question refers to your own personal HPC usage. What is your desired HPC usage for the next 5 years (e.g. CPU/GPU hours)?


Answered: 64 Skipped: 1



| ANSWER CHOICES | RESPONSES | |
|---------------------|-----------|----|
| <1M hours/year | 26.56% | 17 |
| 1M - 5M hours/year | 37.50% | 24 |
| 5M - 10M hours/year | 10.94% | 7 |
| > 10M hours/year | 25.00% | 16 |
| TOTAL | | 64 |

Q7 This question refers to your own personal HPC usage. Where do you prefer to obtain your HPC resources for the next 5 years? If a mixture, please estimate how much of each.

Answered: 62 Skipped: 3



| ANSWER CHOICES | RESPONSES |
|-----------------------|-----------|
| NCI | 45.16% 28 |
| Pawsey | 30.65% 19 |
| OzSTAR | 37.10% 23 |
| Other Domestic | 25.81% 16 |
| Other Overseas | 12.90% 8 |
| Total Respondents: 62 | |

| # | IF A MIXTURE, PLEASE ESTIMATE HOW MUCH OF EACH. | DATE |
|----|--|-------------------|
| 1 | 40% NCI, 45% Pawsey, 15% OzStar | 5/3/2019 12:33 AM |
| 2 | 80% OzSTAR, 20% Other | 4/25/2019 6:32 PM |
| 3 | dedicated Astronomy machine | 4/25/2019 4:11 AM |
| 4 | Probably mostly (80%) other domestic. | 4/18/2019 2:33 PM |
| 5 | NCI 80%, own 20% | 4/18/2019 1:00 PM |
| 6 | 33% each | 4/17/2019 8:48 AM |
| 7 | 50/50 | 4/17/2019 1:58 AM |
| 8 | AWS, Azure | 4/16/2019 4:53 PM |
| 9 | the university | 4/16/2019 3:56 PM |
| 10 | At least 50% domestic: OzSTAR for ad-hoc usage, NCI for specific projects (apply for time) | 4/16/2019 3:47 PM |
| 11 | 80% Ozstar 20% Pawsey | 4/16/2019 3:04 PM |
| 12 | 80% NCI, 20% AVATAR at ANU | 4/16/2019 2:55 PM |
| 13 | NCI: 30%, Kunanyi (University of Tasmania): 70% | 4/16/2019 2:45 PM |

| | | |
|----|---|--------------------|
| 14 | Ideally >80% NCI in the future, but current capacity and CPU hours available means I still require time overseas for the biggest projects. | 4/16/2019 2:43 PM |
| 15 | Open to various ideas/sources. But not interested in spending a lot of time writing multiple proposals. | 4/16/2019 2:36 PM |
| 16 | Of the three machines my preference is OzStar as this is easiest to use, up-to-date and very well managed. However, I would like to try to transition to a more equal share with NCI for particularly intensive jobs. In the past I have used Pawsey, but find the low memory available for jobs to be a hinderance and so would prefer to not use this at this time. | 4/16/2019 1:04 PM |
| 17 | I anticipate that >80% of my resources will be obtained through joint applications with overseas colleagues. | 4/16/2019 12:59 PM |
| 18 | Pawsey as much as possible topping with cloud computing (e.g Azure) if needed | 4/16/2019 12:56 PM |
| 19 | I'm agnostic. Which ever source is easiest and most stable. | 4/16/2019 12:50 PM |
| 20 | 50/50 Ozstar/LIGO data grid | 4/16/2019 12:41 PM |
| 21 | Wherever, as long as it is convenient | 4/16/2019 12:37 PM |

Q8 Any other comments?

Answered: 11 Skipped: 54

| # | RESPONSES | DATE |
|----|--|--------------------|
| 1 | A dedicated investment for astronomy HPC is needed, similar in magnitude to observational hardware and observation time. | 4/25/2019 4:11 AM |
| 2 | Pawsey is easy to access through the partner share, however communication is poor and useful up time is low. | 4/18/2019 3:18 PM |
| 3 | Given my location at Curtin University using Pawsey systems is logistically the easiest solution, and on paper their systems (both magnus and zeus) are suitable for my processing needs. However, there is room for improvement - especially related to file I/O where this is a major bottleneck at many places within our processing pipelines due to the data format we work with. Regularly we receive quicker results on local machines with SSDs for small parts of our pipelines, but the nature of processing 100s of individual datasets means that processing on single local computers is just not possible, and the additional memory, CPU, and disk space generally on the Pawsey systems offers an overall performance improvement. One critique of Pawsey beyond the actual systems themselves is that of lack of communication regarding the astronomy-dedicated scratch space. This partition is regularly problematic, and with a recent update to the firmware (?) causing it to become unusable, we expected better communication from Pawsey that a) there was a problem, and b) what was being done to fix it. As it stands, various support tickets were sent with little helpful communication received for around a week or more after the upgrade. Many CPU-hours were lost because of this lack of communication after this bad upgrade. | 4/17/2019 10:45 AM |
| 4 | HPC is now important even to traditionally observational projects both for image processing and model building. | 4/17/2019 8:48 AM |
| 5 | I am using Pawsey infrastructure because my data is hosted very close to the computers, which minimises data transit. However, the reliability of the supercomputers has not been ideal, resulting in a lot of downtime and making it difficult to progress projects. Communication from Pawsey has not always been good, which compounds the issue. If these problems are not fixed, I will seek resources elsewhere in the long term. I also think this survey is a bit confusing because it asks about "personal" HPC use. I have several students and run a small team; I apply for time for the whole team. Should I count my "personal" use as some fraction of this time? The whole amount? Please bear in mind that different respondents may answer questions differently. I interpreted it to mean "time I applied for / intend to apply for as primary investigator". | 4/16/2019 6:07 PM |
| 6 | I used to be able to make use of national facility time (through the excellent Epic cluster, in Perth). Unfortunately, since that cluster has closed, I have not been able to take advantage of other national facilities, as they are ill-suited to my needs. I did have a sizeable allocation on Raijin three years ago, but was unable to use it due to them being unable to cope with the style of work I do. Specifically, they strongly limited the number of files I could have on the cluster, and after several months of working with them to find a workflow that could allow me to take advantage of the time allocated, we had to abandon the endeavour - and since then, I have stuck to using USQ's Fawkes cluster (as well as taking advantage of the UNSW Katana cluster, in previous years). | 4/16/2019 3:49 PM |
| 7 | Support for high-throughput computing particularly important for gravitational wave data analysis community | 4/16/2019 3:47 PM |
| 8 | To perform world class simulations in my research field (turbulence), we definitely need to get faster and bigger supercomputers to be able to compete with other international groups. | 4/16/2019 3:01 PM |
| 9 | OzSTAR has been incredibly useful from a HPC power point of view, but also hugely valuable is the support, in particular for collaborative access to projects for international researchers. Setting up shared access and user accounts is often very difficult at individual institutions. The OzSTAR team does a fantastic job at this, as well as overseeing a wonderfully easy to use and flexible machine. This is really a very useful resource for the broad community. | 4/16/2019 2:00 PM |
| 10 | I want HPC with a reliable connection (the one to OZSTAR is often going down) and storage. | 4/16/2019 1:31 PM |
| 11 | I am less interested in the processing of large data sets as I am in the storage of large data sets. Once processed in a relatively quick and memory intensive matter I simply use Ozstar to store my data. | 4/16/2019 1:23 PM |