



# ARCHER CSE Service Quarterly Report

Quarter 1 2019



## 1. Document Information and Version History

<b>Version:</b>	1.0
<b>Status</b>	Release
<b>Author(s):</b>	Andy Turner, David Henty, Chris Johnson
<b>Reviewer(s)</b>	Lorna Smith, Alan Simpson

<b>Version</b>	<b>Date</b>	<b>Comments, Changes, Status</b>	<b>Authors, contributors, reviewers</b>
0.1	2019-03-11	Skeleton document	Andy Turner
0.2	2019-04-01	Added Centralised CSE statistics	Andy Turner
0.3	2019-04-02	Added Centralised text	Andy Turner
0.4	2019-04-02	Added eCSE sections	Chris Johnson
0.5	2019-04-02	Added Outreach items	Lorna Smith
0.6	2019-04-02	Added User Survey section	Lorna Smith
0.7	2019-04-02	Added Training section	David Henty
0.8	2019-04-10	Reviewed	Alan Simpson
1.0	2019-04-11	Version for EPSRC	Alan Simpson

## 2. Executive Summary

This report covers the period: 1 January 2019 to 31 March 2019 inclusive.

- Centralised CSE Team:
  - Incorporation of the LASSi parallel I/O performance data into SAFE has allowed us to provide I/O information on all jobs on ARCHER to both users and the Service:
    - Users can use SAFE reporting to gain a better understanding of their I/O use. This information allows them to optimise their use of the system and spot opportunities for software improvement.
    - The Service will gain a better understanding of the I/O use patterns to improve the I/O system configuration and design on current and future national services.
    - The ARCHER CSE team and Cray Centre of Excellence for ARCHER have had a joint paper accepted to the CUG 2019 conference analysing parallel I/O data from LASSi and SAFE, and showing how we can use this data to help users and the service.
  - We have released the second report from the comparative benchmarking effort comparing the single node performance of a variety of UK HPC services and of different processor architectures. The report will help users select the most appropriate HPC resource for their research.
  - In collaboration with the UK Catalyst programme, the ARCHER CSE team has had a paper accepted for the PASC'19 conference comparing performance of Intel processors to Arm processors.
  - We have analysed the feedback from the 2018 ARCHER User Survey and identified a number of opportunities for improvement that we will be implementing in 2019. This analysis can be found in Section 10 below.
- Training:
  - We delivered 9 days (189 student-days) of face-to-face training in the quarter at 4 different locations (with an average feedback score better than “Very Good”) plus 1.0 days of online virtual tutorials with average attendance of 13.
  - To make HPC training as widely accessible as possible, we also ran a 2-day interactive online MPI course over four consecutive Wednesday afternoons with an average of 21 attendees.
  - The ARCHER driving test, which has been running for 4 years, has succeeded in its aim of enabling new users to access the service. From 2015 to 2018, around 350 users gained access to ARCHER. Of these, 275 were active, using around 60% of their total kAU budget.
  - We have had a workshop, *Carpentry for the 99%, not the 1%: Approaches for Hands-on Distance Learning*, accepted for the CarpentryConnect conference in collaboration with the N8 RSE group.
- Outreach:
  - The ARCHER team attended the Big Bang Fair in Birmingham on the 11<sup>th</sup> – 14<sup>th</sup> March again this year, showcasing the value of Supercomputing to the real world, and encouraging children to consider a career in computational science. The event was very busy, with over 80,000 people visiting the fair. Our booth was continuously busy, resulting in a large number of positive interactions with students from a range of locations across the UK.
  - We had a stall at the IET Eng Fest in London on the 2<sup>nd</sup> April 2019, an event organised by the Institution of Engineering and Technology to

provide an “immersive engineering experience” to students. This event was successful with a great deal of interest in the relevance of Supercomputing.

- We provided a workshop at the RBS International Women’s Day event on the 7<sup>th</sup> March. This involved groups of P7 girls learning about STEM activities and we felt this was a successful event to encourage more young women to consider a career in STEM.
- eCSE:
  - All eCSE calls are now complete with 99 of the 100 projects awarded having started and 90 having finished; the final project will be starting early in Q2 2019. We will continue to support running projects to ensure successful delivery of the project work, maximising the benefit to the ARCHER community from the programme funding.

### 3. Collaborations and Outputs Summary

- Presentations:
  - *I/O Usage on ARCHER*, Andy Turner, Storage Challenges in the UK Workshop, 6 Mar 2019, University of Reading
- Meetings:
  - Neelofer Banglawala, Adrian Jackson, *Biomolecular Simulations, Machine Learning and AI HPC Impact Workshop*, 28 January 2019, Oxford
  - Alan Simpson, *Environment HPC Impact Workshop*, 1 February 2019, Manchester
  - Jo Beech-Brandt, *HPC-SIG meeting*, 13 February 2019, London
  - Andy Turner, George Beckett, *Intel Technical Roundtable*, 12-13 March 2019, Edinburgh
  - Andy Turner, Lorna Smith, *RSE Leaders Meeting*, 14 March 2019, Newcastle

## 4. Forward Look

- Centralised CSE Team:
  - We will continue our collaboration with the Cray Centre of Excellence for ARCHER to analyse and understand the I/O use on ARCHER revealed by the Cray LASSi tool and EPCC SAFE to provide benefit to users and the national service.
  - We have an ongoing CSE project to propose machine learning benchmarks for the ARCHER benchmarking activity. We are currently identifying suitable benchmark candidates and in the next quarter we will begin work to evaluate their performance across UK HPC facilities.
  - We are working closely with ARCHER SP to develop a new approach to regression testing of the national HPC service to improve reliability and performance of ARCHER.
  - Based on input from the User Survey, we will look at enhancing the procedures and documentation around centrally-installed software packages.
  - The ARCHER CSE team will be raising the profile of UK HPC at the upcoming CUG 2019, ISC'19 and PASC'19 conferences.
- Training:
  - To ensure that ARCHER users can use a range of future architectures, we are running a Performance Analysis Workshop in Bristol in April that will cover the new ARM64-based Isambard Tier-2 system.
  - With the success of the recent online MPI course, we will also run an online OpenMP course later this calendar year.
- Outreach:
  - Having produced and utilised a new demonstrator for Wee Archie at a series of recent events, we plan to make improvements and enhancements based on input from staff utilising Wee Archie at the events.
- eCSE:
  - As the eCSE programme approaches completion, we will continue to add highlights to the web site, determine programme benefits and create case studies to showcase the success of the programme.

## 5. Contractual Performance Report

This is the contractual performance report for the ARCHER CSE Service for the Reporting Periods: January 2019, February 2019 and March 2019.

The metrics were specified by EPSRC in Schedule 2.2 of the CSE Service Contract.

### CSE Query Metrics

- **QE1:** The percentage of all queries notified to the Contractor by the Help Desk in a Quarter that the Contractor responds to, and agrees a work plan with, the relevant End User within 3 working hours of receiving the notification from the Help Desk. *Service Threshold: 97%; Operating Service Level: 98%.*
- **QE2:** The percentage of all queries notified by the Help Desk to the Contractor that have been satisfactorily resolved or otherwise completed by the Contractor within a 4-month period from the date it was first notified to the Contractor. *Service Threshold: 80%; Operating Service Level: 90%.*
- **TA1:** The percentage of all technical assessments of software proposals provided to the Contractor by the Help Desk in any Service Period that are successfully completed by the Contractor within 10 days of the technical assessment being provided to the Contractor by the Help Desk. *Service Threshold: 85%; Operating Service Level: 90%.*
- **FB1:** The percentage of End User satisfaction surveys for CSE queries carried out in accordance with the Performance Monitoring System by the Contractor showing the level of End User satisfaction to be “satisfactory”, “good” or “excellent”. *Service Threshold: 30%; Operating Service Level: 50%.*

Period Metric	Jan-19		Feb-19		Mar-19		Q1 2019	
	Perf.	SP	Perf.	SP	Perf.	SP	Perf.	Total
QE1	100%	-2	100%	-2	100%	-2	100%	-6
QE2	88%	0	100%	-2	100%	-2	100%	-4
TA1	100%	-1	100%	-1	100%	-1	100%	-3
FB1	100%	-2	100%	-2			100%	-4
<b>Total</b>		-5		-7		-5		-17

*Pink – Below Service Threshold  
Yellow – Below Operating Service Level  
Green – At or above Operating Service Level*

The lower than usual performance for QE2 in January 2019 is due to the closing of a single long-standing In-Depth query where a user requested functionality that was only available in an old version of the HDF5 library. After repeated attempts to contact the user to move the request forward with no response we had to close the query. The user subsequently responded apologising for their lack of response and stating that they no longer required this functionality on ARCHER.

## Training Metrics

- FB2:** The percentage of all training satisfaction surveys carried out in accordance with the Performance Monitoring System by the Contractor) in each Quarter that are rated “good”, “very good” or “excellent”. *Service Threshold: 70%; Operating Service Level: 80%.*

Period	Jan-19		Feb-19		Mar-19		Q1 2019	
	Perf.	SP	Perf.	SP	Perf.	SP	Perf.	Total
FB2	100%	-1	100%	-1	100%	-1	100%	-3
Total		-1		-1		-1		-3

*Pink – Below Service Threshold*  
*Yellow – Below Operating Service Level*  
*Green – At or above Operating Service Level*

## Service Credits

Period	Jan-19	Feb-19	Mar-19
Total Service Points	-6	-8	-6

## 6. CSE Queries

### Queries Resolved in Reporting Period

#### Metric Descriptions

<b>In-Depth</b>	All technical queries passed to ARCHER CSE team
<b>Course Registration</b>	Requests for registration on ARCHER training courses
<b>Course Enquiry</b>	Enquiries about courses
<b>Technical Assessment: &lt;Category&gt;</b>	Request for Technical Assessments of applications for ARCHER time
<b>eCSE Application</b>	Queries relating to eCSE applications

A total of 314 queries were resolved by the CSE service in the reporting period.

<b>Metric</b>	<b>Jan-19</b>	<b>Feb-19</b>	<b>Mar-19</b>	<b>Total</b>
Course Registration	61	111	84	256
eCSE Application	0	0	2	2
In-Depth	8	9	2	19
Course Enquiry	11	4	3	18
Technical Assessment: Grant	3	6	3	12
Technical Assessment: Instant	1	4	2	4
<b>Total</b>	<b>84</b>	<b>134</b>	<b>96</b>	<b>314</b>

2 query feedback responses were received on In-depth queries in the reporting period. This represents a 11% return rate for feedback forms. All responses registered a score of "Excellent". We continue to try to improve the response rate for feedback from queries by offering charity donations for responses and sending additional reminders to users to provide feedback.

Resolved In-Depth queries fell into the following categories:

<b>Category</b>	<b>Number of Queries</b>	<b>% Queries</b>
3rd party software	14	74%
Compilers and system software	1	5%
User programs	1	5%
User behaviour	1	5%
Login, passwords and ssh	1	5%
Performance and scaling	1	5%

## In-Depth Query Highlights

A small number of In-Depth queries have been selected to illustrate the work of the centralised CSE team over the reporting period.

### **Q1149394: Running ICON**

An ARCHER user was trying to compile and run the ICON Earth system modelling application (<https://www.mpimet.mpg.de/en/science/models/icon-esm/>) but was seeing runtime errors with the application. The user was experienced with both ICON and HPC systems and had run the application successfully on another HPC system in Germany but was finding it difficult to get the model running on ARCHER. There were a number of steps to solving this complex query. Firstly, we had to help the user set up their environment on ARCHER correctly for compiling the model; secondly, we had to work with the user to ensure that they were using the most recent release version of the ICON source code; finally, we had to help the user compile the model and test it. It turned out that there were a number of problems causing the issues the user was seeing, including: their environment was not setup correctly to compile the application and the version of ICON that they were using contained a bug that stopped the model from working. The CSE team identified and debugged the issues with the user's environment; then determined the bug in the ICON model and verified that it was fixed in the most recent stable release; and produced a compile process for the user that produced a working version of the model. This allowed the user to continue to use ARCHER for their research. Final feedback from the user:

*I am really happy to say that it looks like ICON is running! Thank you so much for all you help. I really appreciate it.*

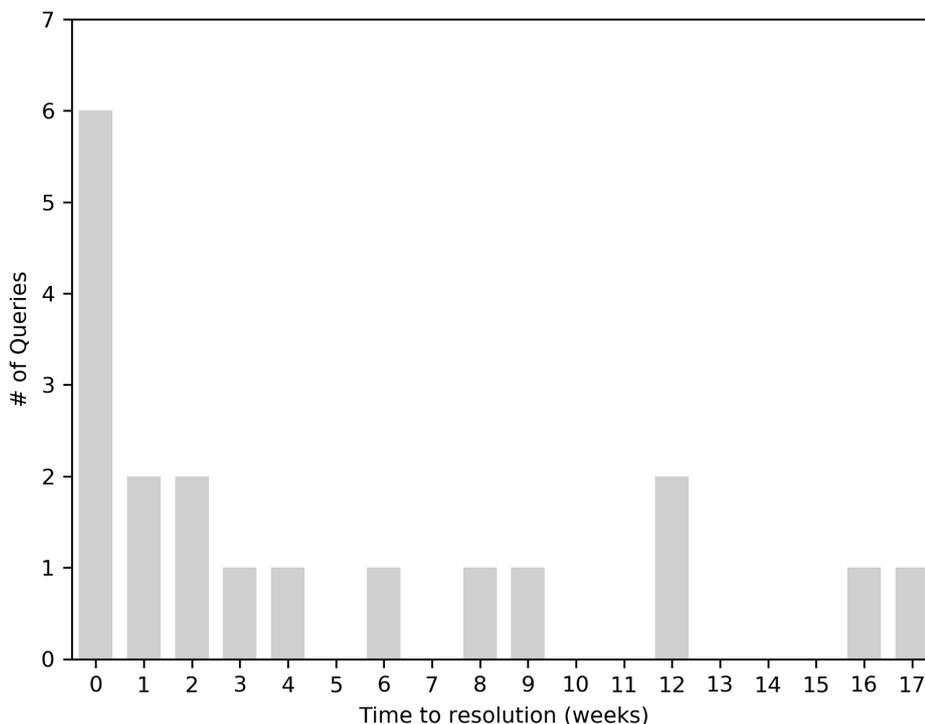
### **Q1171560: Issue with gromacs/2018.2-plumed2.4.2**

A user was seeing issues trying to use the GROMACS (<http://www.gromacs.org>) biomolecular simulation package coupled to the PLUMED free energy calculation framework (<http://www.plumed.org>). They were seeing a strange memory error: `*** glibc detected *** mdrun_mpi: double free or corruption (!prev): 0x000000001719f00 ***`. The CSE team analysed the issue and found that the origin was that the calculation was running out of memory but without triggering an Out Of Memory (OOM) error. We provided the user with advice on how to use the high memory nodes or how to underpopulate compute nodes to make more memory available so they could progress with their research. The user appreciated the help from the CSE team:

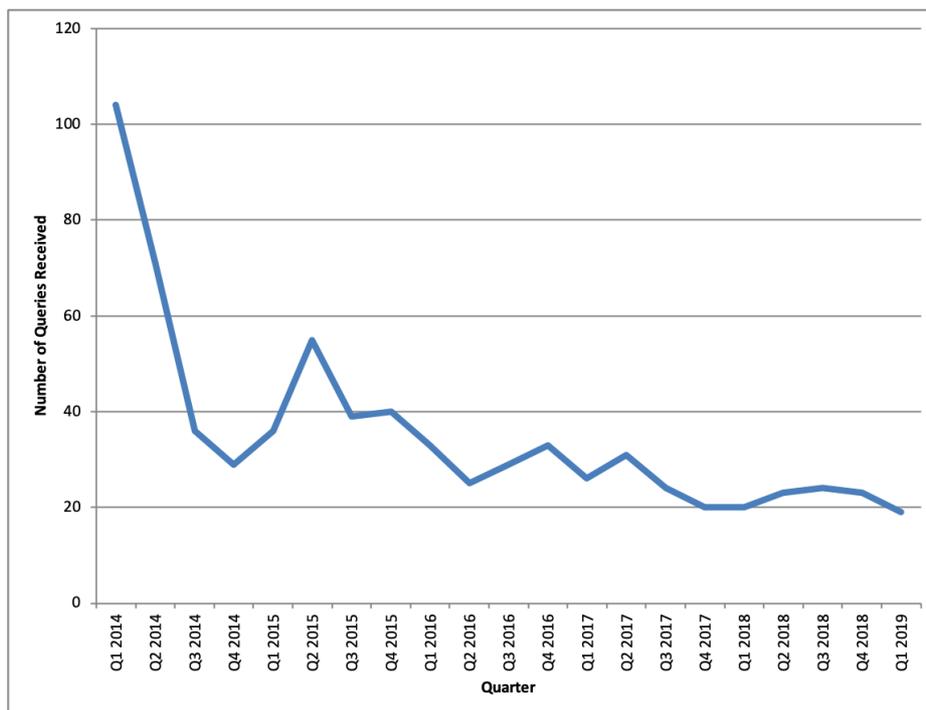
*Mark Bull was very efficient, helpful, and patient - many thanks.*

### In-Depth Query Analysis

The histogram below shows the time to resolution for In-Depth queries in the current reporting period. The median resolution time during this period is 1-2 weeks (median resolution time since 1 Jan 2014 is 5-6 weeks).

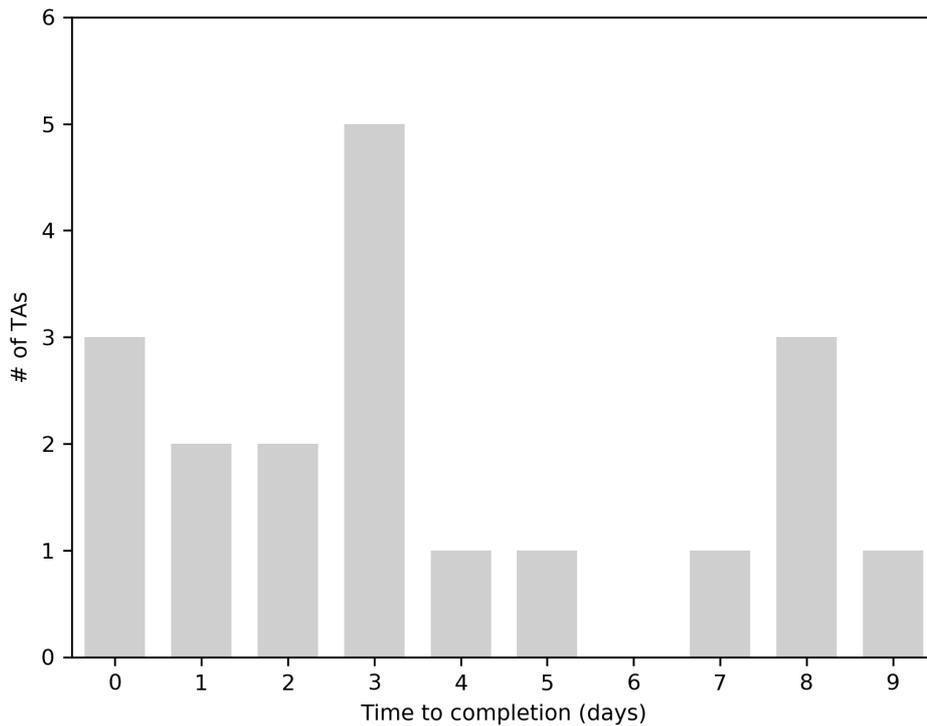


Plot of numbers of In-Depth queries received per quarter:

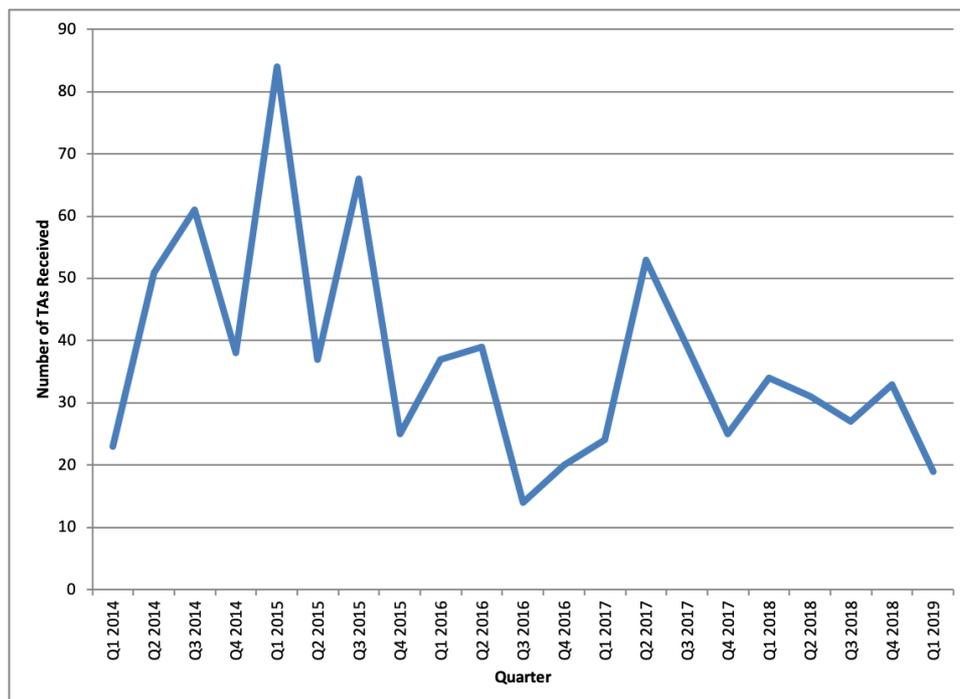


### Technical Assessment Analysis

A histogram of the time to completion for Technical Assessments (see below) reveals that the median completion time for this quarter was 2-3 days (median completion time since 1 Jan 2014 is 2-3 days).



Plot of numbers of Technical Assessments received per quarter:

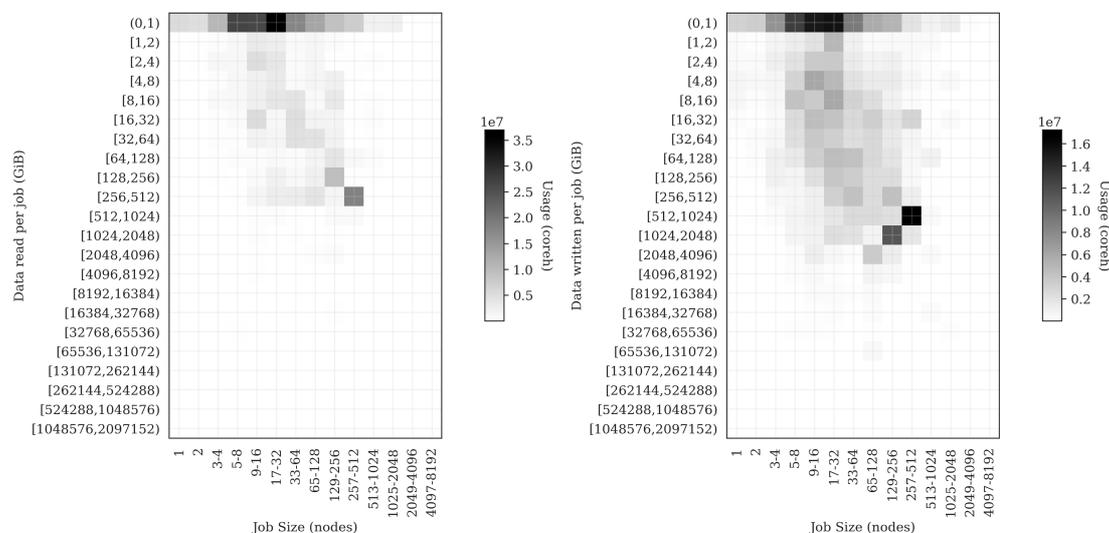


## 7. Centralised CSE Team: Continual Service Improvement

In collaboration with user groups and the other Service partners, the CSE service identified several priority service improvement areas to invest technical effort from the centralised CSE team. This section summarises progress in the reporting period in these areas.

### Parallel I/O: LASSi data in SAFE

As reported in previous quarters, we have configured the import of Cray LASSi I/O statistics into the ARCHER SAFE. This allows us to provide reports for users and application developers that they can use to analyse and understand the I/O usage of their jobs. This information can help users make more efficient use of ARCHER, spot issues with their jobs and help developers identify opportunities to improve their applications. The data in SAFE can also be used by the service to understand I/O patterns across the whole service and look at the different requirements of different research communities and their applications. This information can be used to improve the design of the current service and provide information to help optimise the design of future national services. For example, the heatmap below shows how much data is read and written by jobs of different sizes, weighted by utilisation. We can clearly see a number of different I/O patterns overlaid in this overall view.



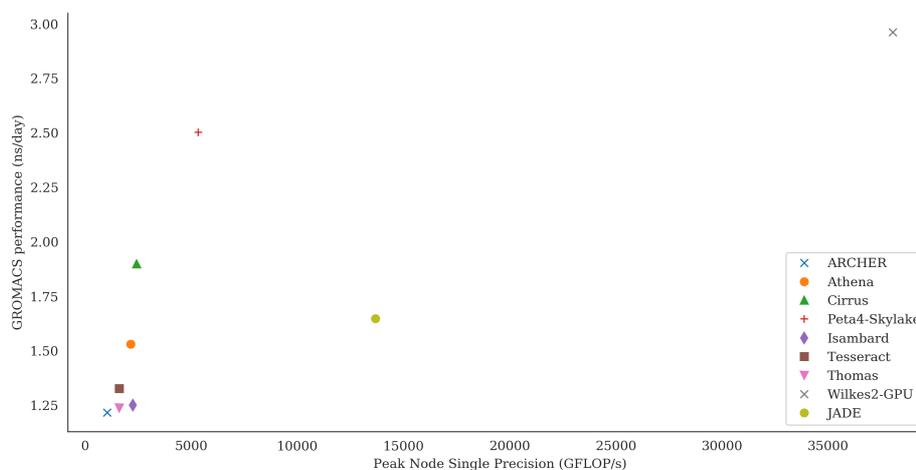
**Heatmap of usage broken down by amount of data read/written within jobs and job size.**

The ARCHER CSE team at EPCC and the Cray Centre of Excellence for ARCHER, along with the I/O expert, Julian Kunkel at the University of Reading, have had a joint paper on this work accepted for the Cray User Group 2019 (CUG 2019) conference from 5-9 May in Montreal, Canada.

### Single node performance comparisons

We have recently published a report comparing and analysing the single node performance for a number of applications across ARCHER, DiRAC and UK Tier-2 HPC systems. To reduce the complexity of the comparisons, we restrict the results in this report to single node only. This allows us to compare the performance of the different compute node architectures without the additional complexity of also comparing different interconnect technologies and topologies. Multi-node comparisons will, however, be the subject of a future report. Architectures compared in this report cover

three generations of Intel Xeon CPUs, Marvell Arm ThunderX2 CPUs and NVidia GPUs. This report is a useful resource to help users select the most appropriate UK HPC resource for their research.



#### Comparison of GROMACS 1400k atom benchmark for a variety of different UK HPC systems.

Following on from this work, the ARCHER CSE team, in collaboration with the UK Catalyst programme<sup>1</sup>, have had a paper accepted for the PASC'19 conference comparing performance of Intel processors to Arm processors.

<sup>1</sup> <https://www.hpe.com/us/en/newsroom/press-release/2018/04/academia-and-industry-collaborate-to-drive-uk-supercomputer-adoption.html>

## 8. Training

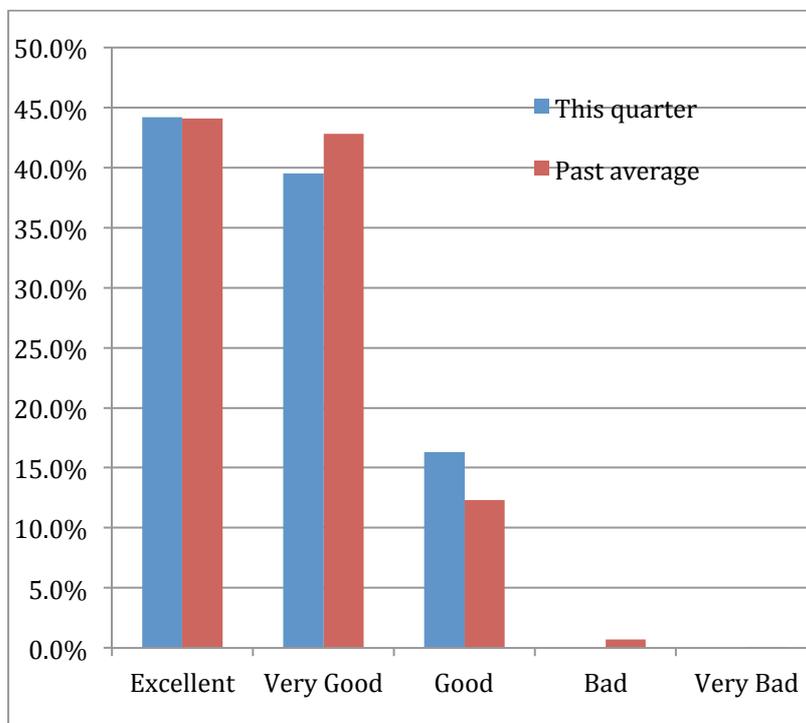
This quarter, the CSE Service has provided a total of 9 days (189 student-days) of face-to-face training across 4 different locations, 2 days of interactive online training (average attendance 21 per session) and 1 day of online tutorials (average attendance 13 per tutorial).

Month	Dates	Course	Location	Days	Attend
Jan 2019	8-9	Data Carpentry	Edinburgh	2	18
	10-11	Introduction to Spark for Data Scientists	Edinburgh	2	29
	17-18	Hands-on Introduction to HPC	Aberdeen	2	15
	21-24	CCP9-MCC-UKCP-EPCC Workshop on Ab initio	Daresbury	1*	25
Feb 2019	5	Periodic Codes	Online	0.5	
	20	Dynamic Load-Balancing in wsiFOAM Online MPI (over 4 consecutive Wednesday afternoons)	Online	2	
Mar 2019	20	Multi-level contact detection in Granular LAMMPS	Online	0.5	
	28-29	Efficient Parallel IO	London	2	20

\*EPCC's role in this collaborative course was to enable attendees to run parallel codes on ARCHER, and we are counting one day as ARCHER training.

We had planned to run a 3-day course on "Introduction to HPC for Life Scientists" in Q1. After discussions with the UK members of the BioExcel project, who are a major target audience, this has been rearranged for Q4 to better fit with their own training timescales.

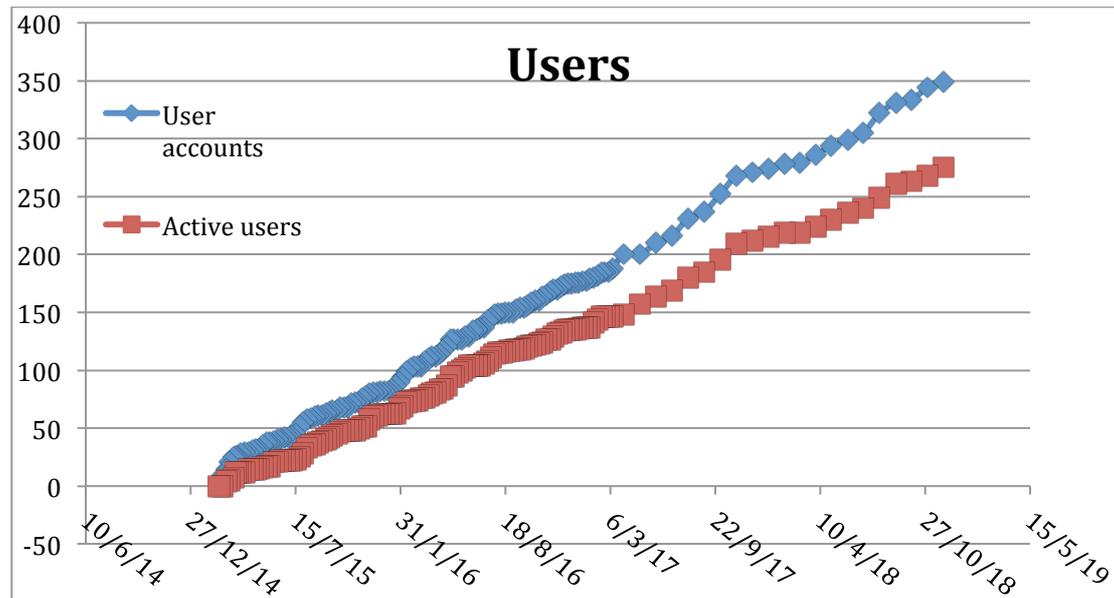
On the feedback for face-to-face courses, attendees rate the course on a scale of 1-5 ("Very bad", "Bad", "Good", "Very good" and "Excellent"). The average feedback using this metric was 4.3, i.e. better than "Very Good". Users provided 43 feedback forms, a response rate of 52%.



22 days of face-to-face training are planned for the first quarter of 2019, plus 1.5 days online.

Month	Dates	Course	Location	Days	Attend
Apr 2019	2-4	Threaded Programming	Southampto	3	
	3-4	GPU Programming with CUDA	n	2	
	16-17	Modern C++ for Computational Scientists	Birmingham	2	
	20	HPC Europa Visitor	Cambridge	0.5	
	24-26	Programme Message-Passing	Online	3	
	24-26	Programming with MPI	Southampto	3	
May 2019	8	Performance Analysis Workshop	n	0.5	
		TBC	Bristol Online		
Jun 2019	19-20	Introduction to Code Saturne	Manchester	2	
	12	TBC	Online	0.5	
	10-11	Software Carpentry	York	2	
	25-27	Advanced OpenMP	Manchester	3	
	TBC	Practical Software Development	Leeds	2	

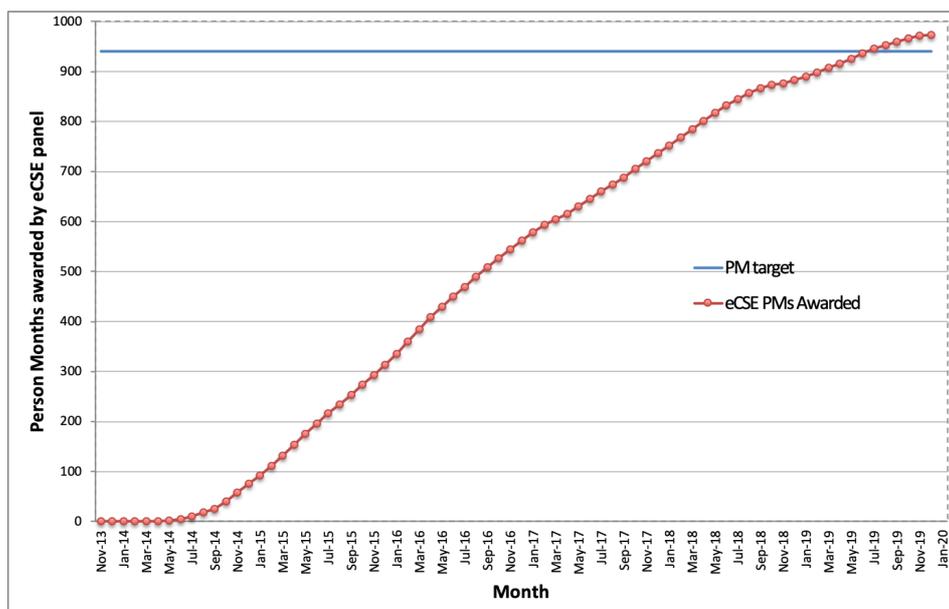
The ARCHER Driving Test has been running since early 2015. The graph below shows the number of users gaining accounts through this scheme, and the number who are active (i.e. have run jobs to ARCHER).



The rate has remained remarkably steady over time, with around 6 new users every month. Users receive 1200 kAUs over a 12 month period; the 275 active users have used a total of just over 200,000 kAUS, some 60% of their total allocation.

It is clear that the driving test is a very effective method of attracting new users to the service, and we believe it should be continued and made available on future national services.

## 9. Embedded CSE (eCSE)



- The eCSE person months awarded up to and including the 13<sup>th</sup> eCSE call are shown in red.
- We committed to awarding at least 941 person months by the end of the project (14 FTEs for 5 years, and 8.4 FTE for year 6).
- 973 person months have been awarded across 100 eCSE projects meaning an extra 32 person months have been awarded over the programme.

eCSE call	No. proposals	No. projects awarded	No. person months awarded	No. projects started	No. projects completed	No. final reports received	Notes
eCSE01	19	14	132	14	14	14	
eCSE02	17	9	82	9	9	9	
eCSE03	16	10	96	10	10	10	
eCSE04	16	8	82	8	8	8	
eCSE05	14	8	94	8	8	8	
eCSE06	9	5	47	5	5	5	
eCSE07	16	5	49	5	5	5	
eCSE08	21	8	88	8	8	6	2 late final reports are being pursued.
eCSE09	19	5	58	5	5	4	1 project completed early due to a staff member leaving and a final report was pursued (see risk list below).
eCSE10	13	6	59	6	6	3	1 late final report is being pursued. 2 final reports are due during the next quarter.
eCSE11	18	6	49	6	6	2	4 late final reports are being pursued.
eCSE12	23	6	41	6	6	5	1 final report is due in the next quarter.
eCSE13	21	10	96	9	0	0	
<b>Total</b>	<b>222</b>	<b>100</b>	<b>973</b>	<b>99</b>	<b>90</b>	<b>79</b>	

- A risk analysis identified all projects as being of either low or very low risk apart from the following which were identified as being of medium risk:
  - eCSE04-10: the PI indicated that the person named to do the technical work may not be available
    - This project went ahead with the original staffing. There was a short delay to the start of the project which started on 01/01/16. The final report was received and reviewed and considered to be a success. This risk will now be removed from the risk list.
  - eCSE08-9: this project had a change of staffing
    - The new staff member was approved by the panel chair and the project has now finished. The final report was received and reviewed and considered to be a success. This risk will now be removed from the risk list.
  - eCSE09-6: this project has terminated early after the recent death of Dr Karl Wilkinson who was one of the Co-Is together with the fact that the researcher doing the work resigned from his current post in Cambridge in November 2018
    - The PI confirmed that the first work package is likely be completed and the project used half its allocated effort. Given the circumstances we agreed to this early termination and the unused funds were used to award eCSE12 projects at the final panel meeting. We have requested a final report to describe the work carried out but this has not yet been received. As the project finished a long time ago we will now assume that the final report for this project will not be received which we believe to be acceptable in the circumstances.
  - eCSE09-8: this project was awarded 19 person months. This is a higher level of effort than awarded for other eCSE projects where 15 person months is the highest level of effort awarded so far
    - Of the 19 months awarded for this project, 7 were for a member of the ARCHER CSE team and the remaining 12 were for an external member of staff at the PI's institution. The project is now complete and the final report is presently under review by the eCSE Panel.
  - eCSE10-5: a change of staffing is required
    - We discussed this with the PI and the project was scaled back and re-staffed. The unused funds were used to fund eCSE12 projects. The project is now complete and we await the final report.
  - eCSE12-20: the project runs right up until 31 October 2018 – almost the end of the CSE contract
    - The project is being monitored via regular contact with the PI.
  - eCSE13-11: the project was awarded partial funding with 2 out of the 9 requested person months being awarded.
    - The PI has confirmed that he agrees to accept the partial funding and a contract has been signed. This risk will now be removed from the risk list.

## 10. ARCHER Annual Survey: CSE Perspective

The ARCHER annual users survey is aimed at understanding the views of users of the current ARCHER service; it seeks feedback looking for areas that are working well and for areas of improvement. The survey measured user satisfaction across a range of areas, with scores from 1 (representing “Very Unsatisfied” to scores of 5 (“Very Satisfied”). The results are shown in table 1.

Service Aspect	2014 Mean Score (out of 5)	2015 Mean Score (out of 5)	2016 Mean Score (out of 5)	2017 Mean Score (out of 5)	2018 Mean Score (out of 5)
<b>Overall Satisfaction</b>	4.4	4.3	4.3	4.4	4.5
<b>Hardware</b>	4.1	4.1	4.2	4.3	3.9
<b>Software</b>	4.0	4.0	4.2	4.1	3.8
<b>Helpdesk</b>	4.5	4.5	4.5	4.6	4.5
<b>Documentation</b>	4.1	4.1	4.2	4.2	4.0
<b>Website</b>	4.1	4.2	4.2	4.2	4.0
<b>Training</b>	4.1	4.1	4.2	4.1	4.3
<b>Webinars</b>	3.6	3.9	3.9	4.2	3.9
<b>Online training</b>	-	4.0	4.1	4.2	3.9

Table 1: Score 5 represents “Very Satisfied”, score 1 “Very Unsatisfied”

188 users responded and overall the service has had very positive feedback. There has been a slight increase in the overall satisfaction rating from 2017, going up from 4.4 to 4.5. Many users provided additional comments, and these are listed in full in the User Survey Report. A few of the ratings have gone down slightly, for example for hardware and software, which perhaps reflects the age of the service and that we are nearing the end of its lifetime. The helpdesk, and access to the CSE expertise via the helpdesk, continues to perform exceptionally well, with an overall rating of 4.5 and a range of positive comments from users.

*Software and Documentation* Looking at the feedback relating to software and documentation, a small number of users commented on centrally-installed software, particularly relating to installed chemistry modelling and simulation software. These comments relate to a desire for more detailed documentation and quicker updates to software packages. We will review the centrally-installed software installation process and the documentation associated with these, focusing particularly on the packages mentioned in the feedback. Documentation, communication and software availability for Python were also mentioned, hence we will review and where appropriate work to enhance these aspects of our Python support.

*Website and Webinars* The overall rating for the website was high, at 4.0. Of the comments received, one looked for improvements for visually impaired people, and we will look to address this comment. A few comments mentioned the desire for better search facilities, and we will also investigate this. Webinars were well received with suggestions for more detailed HPC webinars and case study style webinars highlighting successful publications. We will look to incorporate these suggestions in to the webinar programme.

*Training* Overall satisfaction is again high for training with many positive comments. We will however look to increase the visibility of our online training material and where possible incorporate the positive suggestions for new courses, exercises, certificates, etc.

Finally, there is a small set of comments suggesting additional scripts and tools to support their work. We will review and explore these suggestions and develop any appropriate tools. We will also contact individual users on the survey where their comments suggest they could benefit from CSE support.