



ARCHER CSE Service Quarterly Report

Quarter 2 2019



1. Document Information and Version History

Version:	1.0
Status	Release
Author(s):	Lorna Smith, David Henty, Chris Johnson, Xu Guo
Reviewer(s)	Alan Simpson

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2019-07-01	Skeleton document	Lorna Smith
0.2	2019-07-15	Draft of the report	Lorna Smith
0.3	2019-07-15	Added eCSE section	Chris Johnson
0.4	2019-07-15	Added training section	David Henty
0.5	2019-07-15	Added CSI section	Xu Guo
0.6	2019-07-16	First full draft of the report	Lorna Smith
0.7	2019-07-16	Reviewed	Alan Simpson
1.0	2019-07-17	Version for EPSRC	Alan Simpson, Lorna Smith

2. Executive Summary

This report covers the period: 1 April 2019 to 31 June 2019 inclusive.

- Centralised CSE Team:
 - The CSE service and the Cray Centre of Excellence presented a joint paper at the Cray User Group meeting in Montreal from 5-9 May 2019 describing our collaborative work on understanding parallel I/O performance. The paper was well received and showcased ARCHER's and UK HPC's impact on an international stage. The outputs from this work continue to be used in day-to-day operations on ARCHER.
 - The ARCHER CSE and SP teams have been working closely to prepare and support users in making the most effective use of storage as the end of service approaches. The Data Management guide is being updated and a webinar is planned for September.
 - Based on feedback from the Annual Survey, a review of the significant package accounts has been undertaken to ensure they are up-to-date.
- eCSE:
 - All eCSE calls are now complete with 91 of the 100 awarded projects now finished, and the remaining 9 due to finish by the end of November 2019. 84 of the projects have submitted final reports. During the period we have focused on supporting on-going projects and on maximising impact and benefit from completed projects.
- Training:
 - We delivered 21 days (432 student-days) of face-to-face training in the quarter at 7 different locations. The average feedback score was better than "Very Good", showing our on-going commitment to high-quality training throughout the entire service lifetime.
 - Two of the courses focused on specific software packages on ARCHER: Code Saturne and UKAMOR. In addition, we assisted at training courses on the Unified Model and on CASTEP.
 - We co-organised a workshop at "Carpentries Connect" in Manchester on 26 June 2019 titled "Approaches for Hands-on Distance Learning" which discussed experiences of distance learning with the RSE community, aiming to distil elements of best practice.
- Outreach:
 - The ARCHER team presented Wee Archie at the Royal Highland Show (20-23 June). We were invited by the Department for BEIS and were stationed in the UK Government tent. We raised awareness of the ARCHER service to members of the public. David Mundell, Secretary of State for Scotland was among the people who used Wee Archie.
 - An HPC Champions event was held on the 9th May in Edinburgh (alongside the WHPC launch – see following). Sessions included site updates from all Tier-2 centres, the current HPC Landscape which provided talks on the UKRI roadmap, the transition between ARCHER and ARCHER2 and a lightning talk session. The day brought together a wide range of HPC champions sharing best practice and experience.
 - We were very pleased to be accepted as one of the first WHPC chapters, developing a chapter that encompasses all of the UK. We held a successful launch event on the 8th May, with attendees from different organisation from across the UK. The chapter has a strong focus on diversity and we were keen to use the launch to show that positive cultural changes in the workplace can benefit everybody.

3. Collaborations and Outputs Summary

- Presentations:
 - *ARCHER: HPC in the UK and beyond*, **Chris Johnson**, Presentation given at HPC Day, 8 April 2019, Coventry University
 - *ARCHER and the eCSE programme*, **Chris Johnson**, Presentation given to RSE group, 24 June 2019, University of Leeds
 - *Carpentry for the 99%, not the 1%: approaches for hands-on distance learning*, **Andy Turner and Martin Callaghan**, 25-17 June 2019, Manchester, UK
 - *ARCHER and EPCC's support of WHPC*, **Neelofer Banglawala**, WHPC reception at ISC, 19th June 2019, Frankfurt, Germany
- Exhibitions and Outreach Events:
 - *ISC*, Exhibition booth promoting ARCHER to exhibition visitors using Wee Archie, **Neelofer Banglawala**, Sunday 16th - Thursday 20th July 2019, Frankfurt, Germany
 - *Royal Highland Show*, Stationed in the UK Government tent we raised awareness of the ARCHER service to members of the public and to politicians, **David Henty, Andy Turner, Jo Beech-Brandt, Lorna Smith, Sypro Nita, Alistair Grant**, 20 June -23 June, Edinburgh, UK
 - *IET Engfest 3.0*, Wee Archie demonstration to school children, **Juan Rodriguez Herrera, Ben Morse**, 2 April, IET Savoy Place, London, UK
- Meetings:
 - *CP2K Users' and Developers' meeting*, **Arno Proeme**, June 14 June 2019, Imperial College, Molecular Sciences Research Hub, London, UK
 - *ARCHER SAC*, **Lorna Smith and Alan Simpson**, 10th April 2019, London UK
 - *HEC Consortia Meeting*, **Lorna Smith**, 22nd May 2019, London UK
 - *ARCHER RAP Meeting*, **Andy Turner**, Technical Advisor, 19 June 2019, Swindon, UK
 - *Scottish Regional HPC Technical Meeting*, **Andy Turner**, 7 June 2019, Dundee, UK
 - *Society of Research Software Engineering, Trustees Meeting*, **Andy Turner**, 8-9 May 2019, London, UK
 - *EPSRC ARCHER and Tier2 RAP pre-panel meeting*, **Andy Turner**, 3 May 2019, Online
 - *ARCHER Champions*, **Lorna Smith, Alan Simpson, Jo Beech-Brandt, Clair Barrass**, 9 May 2019, Edinburgh, UK
 - *EPCC's Women in HPC Launch*, **Lorna Smith, Neelofer Banglawala, Alan Simpson**, 8 May 2019, Edinburgh, UK
- Papers
 - *Evaluating the Arm Ecosystem for High Performance Computing*, **Adrian Jackson, Andy Turner, Michele Weiland, Nick Johnson, Olly Perks, Mark Parsons**, PASC19, 12 June -14 June 2019, Zurich, Switzerland.
<https://doi.org/10.1145/3324989.3325722>

4. Forward Look

- Centralised CSE Team:
 - We are looking at ways of ensuring we support users as the end of service approaches, for example around how best to manage their data. Current efforts are focused on updating the Data Management Guide and a new release is planned for the end of July.
 - We plan to continue to enhance our support for user communities in the use of the larger, community based application codes. In addition to the work on package accounts, a number of training courses have had an application focus. We will be holding a training course associated with an advanced VAMPIRE workshop in July.
 - For the parallel I/O work, we plan to include statistics on metadata operations in SAFE reporting, to allow analysis of this aspect of parallel I/O performance.
 - During this period we plan to use the ARCHER benchmarks to test the functionality and performance of Singularity containers for different HPC use cases. The evaluation of performance and convenience of using containerised HPC applications will be shared with ARCHER/Cirrus users via an ARCHER white paper and an ARCHER webinar.
- Training:
 - To assist users with their data management towards the end of the ARCHER service, we are running a webinar in September.
 - With the success of our new Data Science course on Spark, first run in Q1 of this year, we plan to run this again in 2019 and are investigating using the Alan Turing Institute in London as a venue.
- eCSE:
 - We will continue to focus on ensuring the maximum impact is delivered from the eCSE programme, through scientific highlights, case studies and benefits realisation.
 - We will also continue to support on-going projects to ensure successful delivery of the technical work, including final reports and impact/benefits data, maximising the benefit to the ARCHER community from the programme funding.
- Outreach
 - Our recently developed costal erosion / flood defence demonstration has proved very successful at recent outreach events. Over the next period we will look at further enhancements to this demonstration, based on feedback from event attendees.

5. Contractual Performance Report

This is the contractual performance report for the ARCHER CSE Service for the Reporting Periods: April 2019, May 2019 and June 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	Apr-19		May-19		June-19		Q2 2019	
	Perf.	SP	Perf.	SP	Perf.	SP	Perf.	Total
QE1	100%	-2	100%	-2	100%	-2	100%	-6
QE2	100%	-2	100%	-2	100%	-2	100%	-6
TA1	100%	-1	100%	-1	100%	-1	100%	-3
FB1	100%	-2	100%	-2	100%	-2	100%	-6
Total		-7		-7		-7		-21

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	Apr-19		May-19		June-19		Q2 2019	
	Perf.	SP	Perf.	SP	Perf.	SP	Perf.	Total
FB2	98%	-1	100%	-1	100%	-1	99%	-3
Total		-1		-1		-1		-3

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

The 98% in April 19 relates to one survey rated as “bad” from 18 returned surveys for a course on Modern C++ for Computational Scientists. The remaining 17 surveys were rated between “good” to “excellent”.

Service Credits

Period	Apr-19	May-19	June-19
Total Service Points	-8	-8	-8

6. CSE Queries

Queries Resolved in Reporting Period

Metric Descriptions

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

A total of 383 queries were resolved by the CSE service in the reporting period.

Metric	Mar-19	Apr-19	May-19	Total
Course Registration	137	122	61	320
In-depth	10	9	5	24
Course Enquiry	3	4	3	11
Technical Assessment: Grant	1	5	2	8
Technical Assessment: Instant	6	3	1	10
Technical Assessment: RAP	9	1	0	10
Total	166	144	72	383

6 query feedback responses were received on In-depth queries in the reporting period. This represents a 25% return rate for feedback forms. Five responses registered a score of "Excellent" and one of "Good". 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:

Category	Number of Queries	% Queries
3rd party software	15	62.5%
Compiler and system software	3	12.5%
Batch system and queues	2	8.3%
User behaviour	1	4.2%
Disk, tape, resources	1	4.2%
Access to HPC	1	4.2%
User programs	1	4.2%

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.

Q1204064: Code not running

An ARCHER user found that his in-house CFD code (Hydro3D) was hanging on ARCHER, but running on the Supercomputing Wales system. The user believed the problem had occurred two years ago on ARCHER and been fixed at the time. We were able to get the code running successfully on ARCHER and explain to the user why it ran on Supercomputing Wales and not on ARCHER initially. We identified two issues. The first related to a tag space problem and by refactoring the code to make the tag offsets parameters and reducing the values of these, we were able to fit within the Cray MPI tag limit. We also identified a deadlock occurring in the MPI_SEND calls, a problem we resolved by increasing the maximum size of the non-blocking messages to 128KB. We explained to the users that the Intel MPI on the Supercomputing Wales system had a higher non-blocking message size by default than Cray MPI, explaining the difference in behaviour on the two systems.

I'm extremely satisfied with the support that Mark has given me during this issue. A very well done job that not only sorted our problem but improved the speed up of our code.

Q1216836: Help to compile and install the Qb@ll package on archer

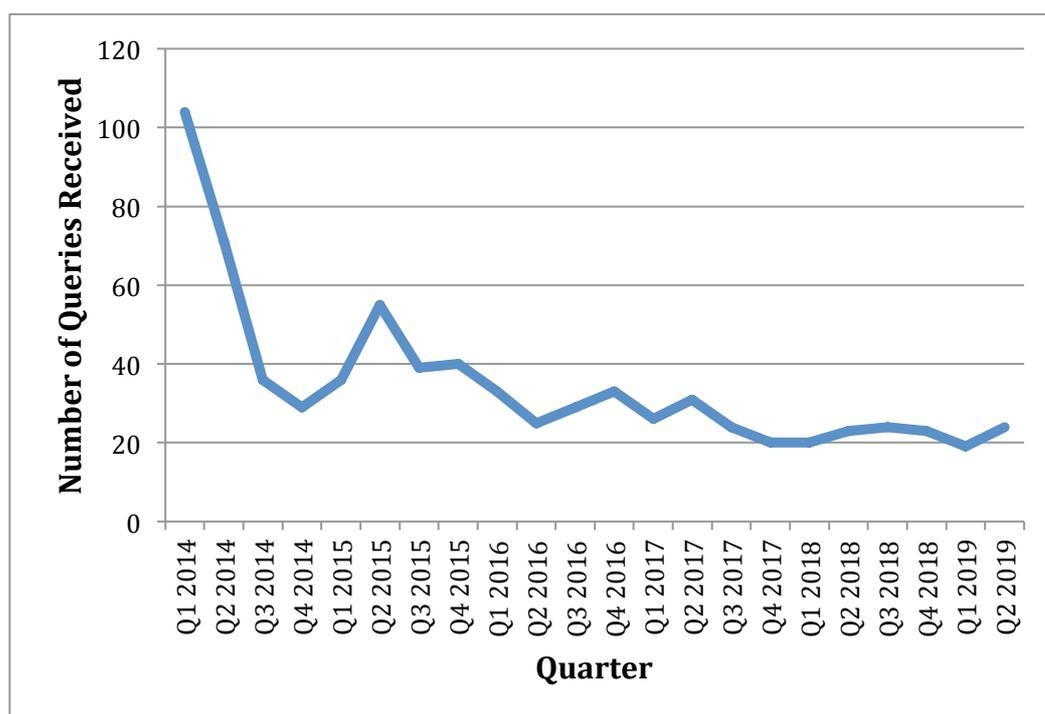
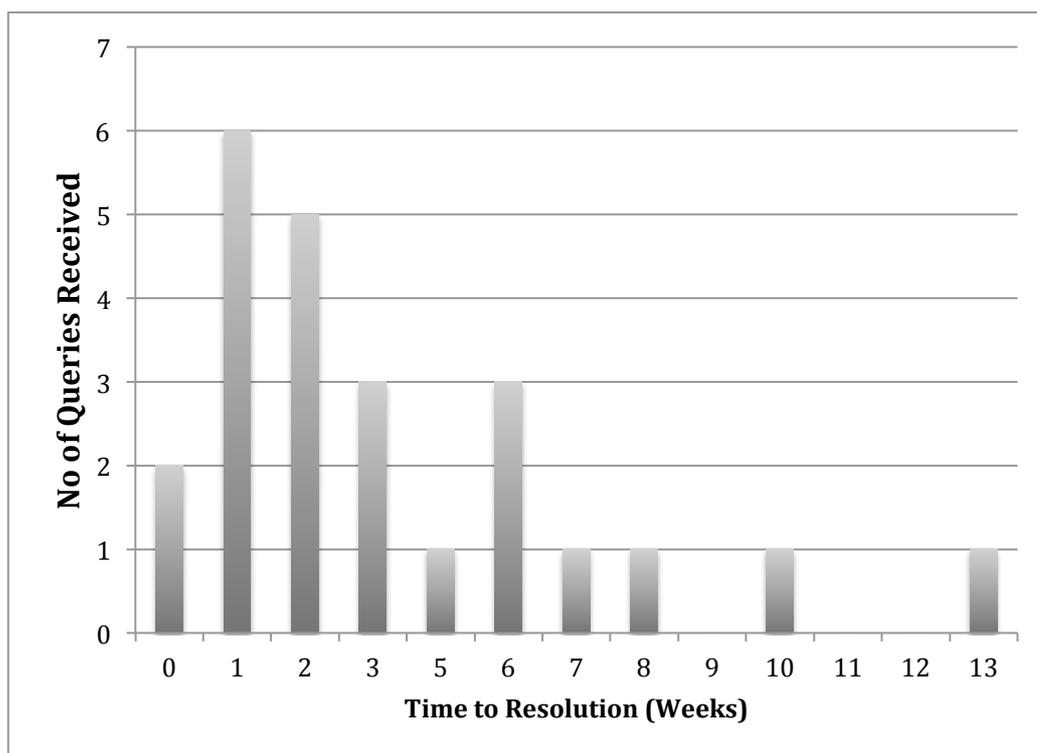
A user requested help to install Qb@ll on ARCHER. We were able to install the package for the user and provide them with instructions on how to install the package themselves. Installation was relatively straight-forward but did involve utilisation of the Intel MKL Link Line to show how to set the LIBS_BLAS environment variable. The user was able to install and test Qb@all successfully.

Thanks a lot for so many help! I have installed qball following the steps as suggested. And the test passed also with the same results! 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 2-3 weeks.

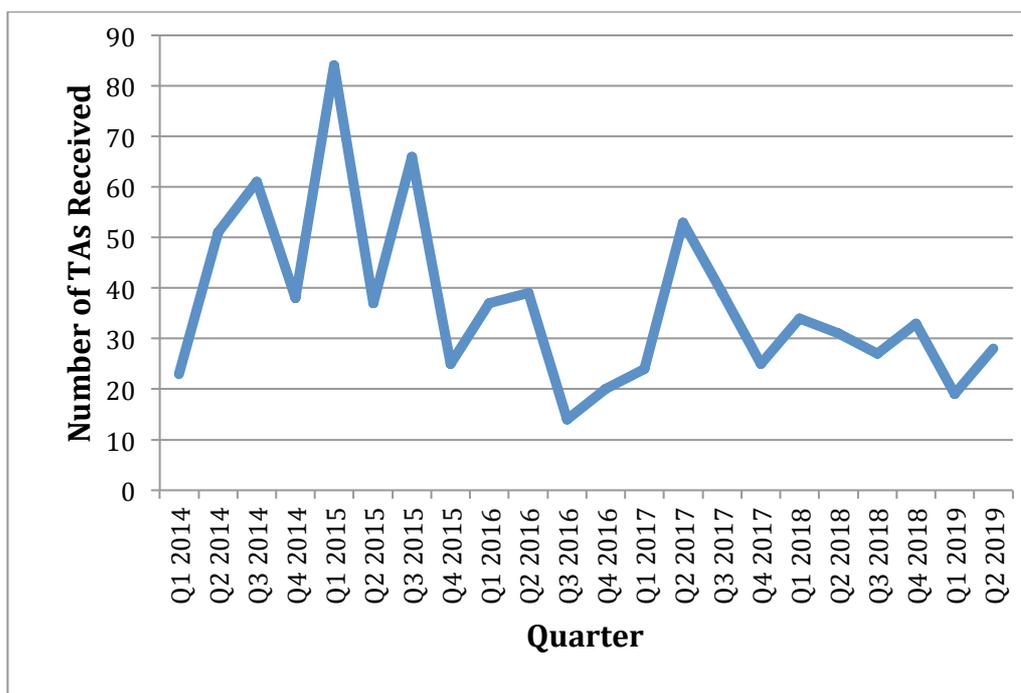
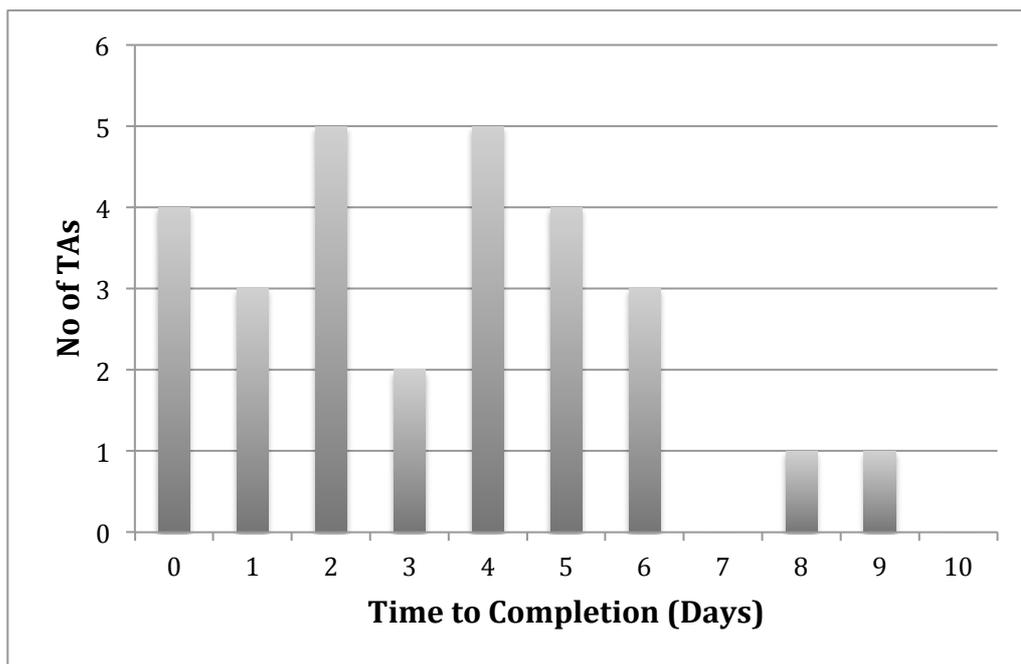
Plot of numbers of In-Depth queries received per quarter:



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 3-4 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 provides highlights from the reporting period.

Parallel I/O Performance Variation

The ARCHER CSE service and the Cray ARCHER Centre of Excellence presented a joint paper at the Cray User Group meeting in Montreal from 5-9 May 2019 describing our collaborative work on understanding parallel I/O performance and use on ARCHER using the Cray LASSi and EPCC SAFE tools. The paper was well-received, stimulating interesting discussion and showcasing ARCHER's and UK HPC's impact on an international stage. The outputs from this work continue to be used in day-to-day operations on ARCHER to identify parallel I/O issues affecting the performance of applications on ARCHER. Users have access to I/O stats through the SAFE. The next stages in this project are: to include statistics on metadata operations in SAFE reporting, to allow analysis of this aspect of parallel I/O performance; and to continue to compare I/O performance data with other HPC systems around the world.

Tier-2 CPU benchmarking

The open source nature of this benchmarking project has allowed external input and suggestions to enhance the activity. For example, recent input around discrepancies in clock speed for different SIMD instructions compared to base processor frequencies has improved the understanding of performance for a number of different benchmarks. Work on this project continues with the inclusion of results from the HPE ARM-based Catalyst system at EPCC and results from the latest, Cascade Lake generation of Intel Xeon processors. Two of the Tier-2 centres (MMM Hub and HPC Mid+) plan to contribute additional results over the next few months to expand the amount of data in the repository.

Containers for HPC

Following on from our previous report, we have focused on using Singularity containers (<https://www.sylabs.io/>). We plan to use the ARCHER benchmarks to test the functionality and performance of Singularity containers for different HPC use cases and have started to benchmark the synthetic HPPC benchmark to compare its "bare metal" performance against that achieved within a Singularity container running on Cirrus. The evaluation of performance and convenience of using containerised HPC applications will be shared with ARCHER/Cirrus users via an ARCHER white paper and an ARCHER webinar.

ARCHER Data Management Guide update

Working closely with the SP team, we have been looking at ways to help prepare and support users in making the most effective use of storage as the end of service approaches. The Data Management guide is being updated, focusing on data transfer protocols, data management plans and the different available file systems. A new release is planned for the end of July with a webinar planned for September. We anticipate running multiple such webinars before the end of service

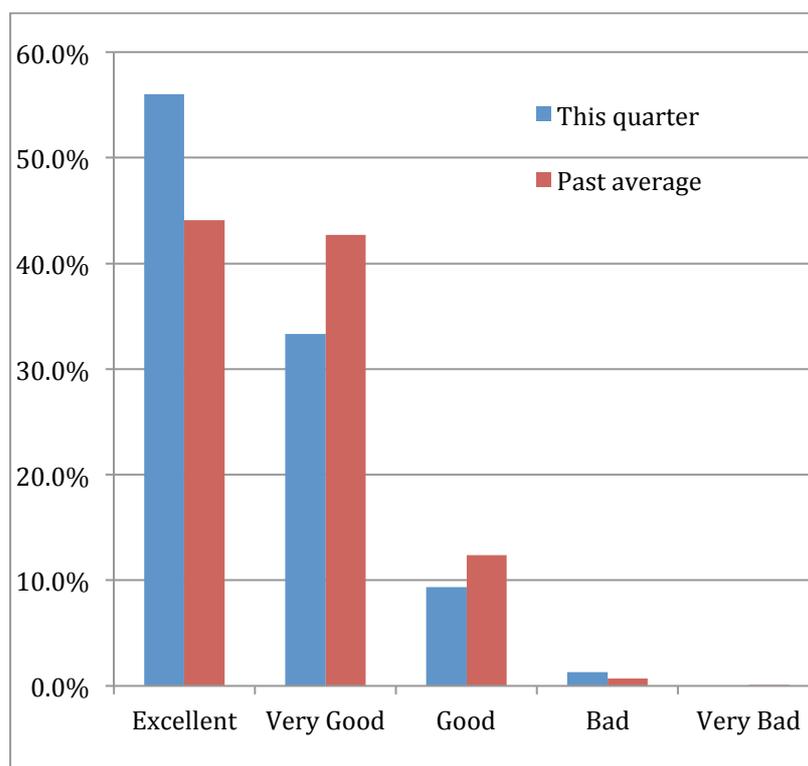
8. Training

This quarter, the CSE Service has provided a total of 21 days (432 student-days) of face-to-face training across 7 different locations and 1.5 days of online tutorials (average attendance 14 per tutorial).

Month	Dates	Course	Location	Days	Attend
Apr 2019	2-4	Threaded Programming	Southampton	3	33
	3-4	GPU Programming with CUDA	Birmingham	2	17
	16-17	Modern C++ for Computational Scientists	Cambridge	2	25
	24	HPC Europa Visitor Programme	Online	0.5	
	24-26	Message-Passing Programming with MPI	Southampton	3	25
	24-26	Performance Analysis Workshop	Bristol	3	14
May 2019	22	Open Source HPC Benchmarking	Online	0.5	
Jun 2019	19-20	Introduction to Code Saturne	Manchester	2	17
	10-11	Software Carpentry	York	2	22
	24-26	R-MADAM 2019 Workshop	London	1*	30
	25-27	Advanced OpenMP	Manchester	3	8
	26	Modern Fortran	Online	0.5	

*This was a 3-day workshop organised by the UK-AMOR consortium. As part of this, the CSE team provided one day of hands-on training on how to use their R-matrix software on ARCHER.

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 4.4, i.e. better than "Very Good". Users provided 75 feedback forms, a response rate of 47%.



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

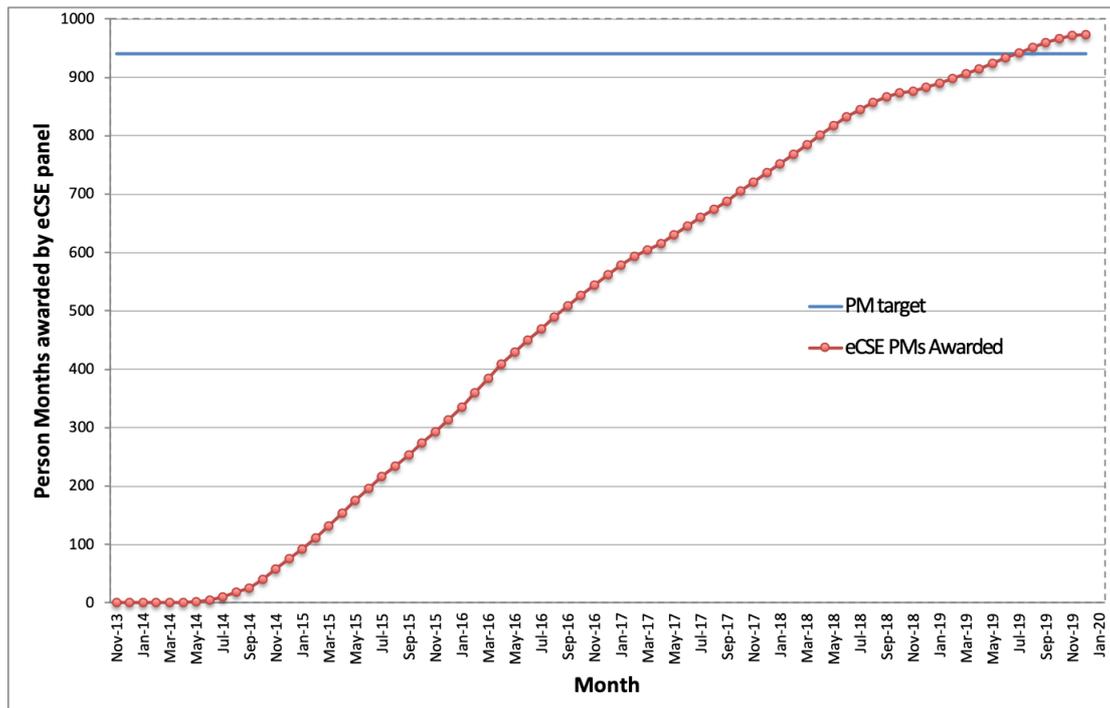
Month	Dates	Course	Location	Days	Attend
Jul 2019	4-5	GPU Programming with CUDA	London	2	
	10-11	Practical Software Development	Leeds	2	
	15-16	Hands-on Introduction to HPC	Edinburgh	2	
	17-19	Message-Passing Programming with MPI	Edinburgh	3	
	22-23	Advanced MPI	Oxford	2	
	22-26	2 nd Advanced VAMPIRE Workshop	York	1*	
	30-31	Single Node Performance Optimisation	London	2	
Aug 2019	TBD	Performance of Reproducible Computations	Online	0.5	
	7-9	Parallel Design Patterns	Durham	3	
	14	OpenMP 5.0	Online	0.5	
Sep 2019	27-28	Modern C++ for Computational Scientists	Glasgow	2	
	11	Data Transfer from ARCHER	Online	0.5	
	Sep 30	Porting and Optimisation	Cambridge	2	
	- Oct 1	Workshop			

*As part of this weeklong consortium workshop, the CSE team is providing one day of hands-on training on how to run the software on ARCHER.

9. Embedded CSE (eCSE)

All calls for the eCSE programme have now finished and all the effort has been awarded. The section below summaries the number of person months awarded over time, highlighting that the service was able to deliver more person months than required contractually.

Overview of eCSE effort



- The eCSE person months awarded up to and including the 13th 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 FTEs 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.

Project Summary

The table below summaries the status of all the projects funded across each eCSE call. 91 of a total of 100 projects have now completed, with 84 of these having submitted final reports. We continue to monitor and work with groups to ensure successful completion of projects and to ensure impact summaries are publicised on the web.

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	8	
eCSE09	19	5	58	5	5	4	Due to extenuating circumstances reported previously, we will not receive a report for project eCSE09-6.
eCSE10	13	6	59	6	6	4	2 final reports are due during the next quarter.
eCSE11	18	6	49	6	6	4	2 late final reports are being pursued.
eCSE12	23	6	41	6	6	5	1 late final report is being pursued.
eCSE13	21	10	96	10	1	0	1 final report is due in the next quarter
Total	222	100	973	100	91	84	

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:

- 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 final report is presently under review by the eCSE Panel.
- eCSE10-5: a change of staffing was 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.