



ARCHER CSE Service Quarterly Report

Quarter 3 2019



1. Document Information and Version History

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

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2019-07-10	Skeleton document	Lorna Smith
0.2	2019-10-10	Complete first draft	Lorna Smith, Chris Johnson, Xu Guo, David Henty
0.3	2019-11-10	Reviewed	Alan Simpson
1.0	2019-11-14	Version for EPSRC	Lorna Smith, Alan Simpson

2. Executive Summary

This report covers the period: 1 July 2019 to 30 September 2019 inclusive.

- Centralised CSE Team:
 - We have been investigating the new I/O module of OpenFOAM, to help inform and provide guidance to existing users. There have been problems in the past with the number of files created by OpenFOAM, so this should be beneficial to the community.
 - We have recently updated the ARCHER Data Transfer Guide and ran a virtual tutorial on “Good Practice for Transferring Data”. This is to help prepare users for the end of the ARCHER service.
 - The ARCHER Image Competition is currently open and has already had a strong set of entries.
- eCSE:
 - All eCSE calls are now complete with all 92 of the 100 awarded projects having finished; the remaining 8 are in progress and due to finish by the end of the ARCHER Service. 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.
 - 82 scientific highlights are now available on the website, and two of these are provided below to showcase the science.
- Training:
 - We delivered 21 days (378 student-days) of face-to-face training in the quarter at 8 different locations (with an average feedback score better than “Very Good”), plus 1.5 days of online virtual tutorials with average attendance of 15.
 - We continue to train users in a range of upcoming HPC technologies including new ways to program GPUs and a collaborative course with ARM on their new ARM64 processor.
 - Our new course on “C++ for Computational Scientists” continues to be very popular, with 30 attendees at the run in August.
- Outreach and Engagement:
 - HPC Champions was held in Birmingham in September; this was co-located with RSECon19. There was a wide variety of talks and a discussion session on how we can build a UK HPC knowledge base to better share expertise between ourselves and to the user communities.
 - We had a booth at The Royal Highland Show in Edinburgh in June, as part of the UK Government tent. This provided a diverse audience and we were able to highlight the relevance and importance of HPC to the general public.

3. Collaborations and Outputs Summary

- Presentations:
 - “HPC Carpentry”, HPC Champions, Andy Turner, 16 Sep 2019, University of Birmingham, UK

- Meetings:
 - RSECon19, Jo Beech-Brandt, Andy Turner, Birmingham, UK
 - HPC Champions, Jo Beech-Brandt, Andy Turner, 16th September, Birmingham, UK
 - HPC-SIG Meeting, Andy Turner, 1 Jul 2019, Queens University Belfast, UK
 - DiRAC Day, Andy Turner, 11-12 Sep 2019, University of Leicester, UK

- Exhibitions and Outreach Events:
 - Royal Highland Show, Lorna Smith, Spyro Nita, Andy Turner, David Henty, Jane Kennedy, Chris Wood, Gordon Gibb, 2-23rd June, Ingliston, Edinburgh
 - Bang Goes the Borders, Lorna Smith, Chris Wood, Jane Kennedy, Mario Antonioletti, 21st September, Melrose.

4. Forward Look

- Centralised CSE Team:
 - As we approach the end of the ARCHER service, we will continue to support users in the transition to other HPC systems, for example through best practice around data management.
- Training:
 - We are running a 1-day online course on the LAMMPS package over two consecutive Wednesday afternoons, and giving attendees access to ARCHER for the practical exercises. This course has already attracted over 30 registrations.
 - The same format is being used for a two-day OpenMP course, also run online to maximise accessibility. Both courses will be recorded and videos made available on the ARCHER web pages for users who are unable to attend any sessions.
 - Containers are becoming ever more popular as a means to ensure portability of software between different systems. In response to this we have developed a new 1-day course on “Containers for HPC” to be run in the coming quarter.
 - The content of the new course “Modern C++ for Computational Scientists” will be reviewed based on user feedback regarding the programming exercises.
 - The “Introduction to Spark for Data Scientists” course is being run in collaboration with The Alan Turing Institute at their headquarters in London.
- eCSE:
 - The majority of the eCSE projects are now finished and our focus is, in addition to supporting the remaining projects, on demonstrating the impact and benefit of the programme through both qualitative and quantitative outputs.
- Outreach
 - The outreach team has a booth at New Scientist Live in London in October. The audience includes school children and the general public and is an ideal opportunity to showcase the relevance of Supercomputing to society.
 - We continue to support work experience placements for secondary school children, providing week-long placements. We anticipate hosting two students during this quarter.

5. Contractual Performance Report

This is the contractual performance report for the ARCHER CSE Service for the Reporting Periods: July 2019, August 2019 and September 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	July-19		Aug-19		Sept-19		Q3 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	July-19		Aug-19		Sept-19		Q2 2019	
	Perf.	SP	Perf.	SP	Perf.	SP	Perf.	Total
FB2	100%	-1	92%	-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 92% in August relates to one survey rated as “bad” from 8 returned surveys for a course on Modern C++ for Computational Scientists. The remaining 7 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 267 queries were resolved by the CSE service in the reporting period.

Metric	Mar-19	Apr-19	May-19	Total
Course Registration	58	33	124	215
In-depth	7	11	4	22
Course Enquiry	4	5	5	14
Technical Assessment: Grant	6	3	2	11
Technical Assessment: Instant	1	0	1	2
Technical Assessment: RAP	0	0	1	1
Technical Assessment: HEC	0	0	1	1
eCSE	0	1	0	1
Total	76	53	138	267

2 query feedback responses were received on In-depth queries in the reporting period. This represents a 9% return rate for feedback forms. Both 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:

Category	Number of Queries	% Queries
3rd party software	14	63.6%
Compiler and system software	4	18.2%
Batch system and queues	1	8.3%
Login, password and ssh	1	4.5%
Node failure	1	4.5%
User programs	1	4.5%

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.

Q1228040 Help with VASP

A user had encountered a problem running VASP when performing phonon calculations. There was a known solution to this, however the user wanted assistance in creating a new version of VASP that incorporated this solution. David was able to create a new version of VASP for the user incorporating these changes. The user was then able to run their simulations. This type of support is representative of a number of our in-depth queries, involving helping users to utilise these large, community codes.

“David helped to compile a new version of vasp for me only. I appreciate it very much.”

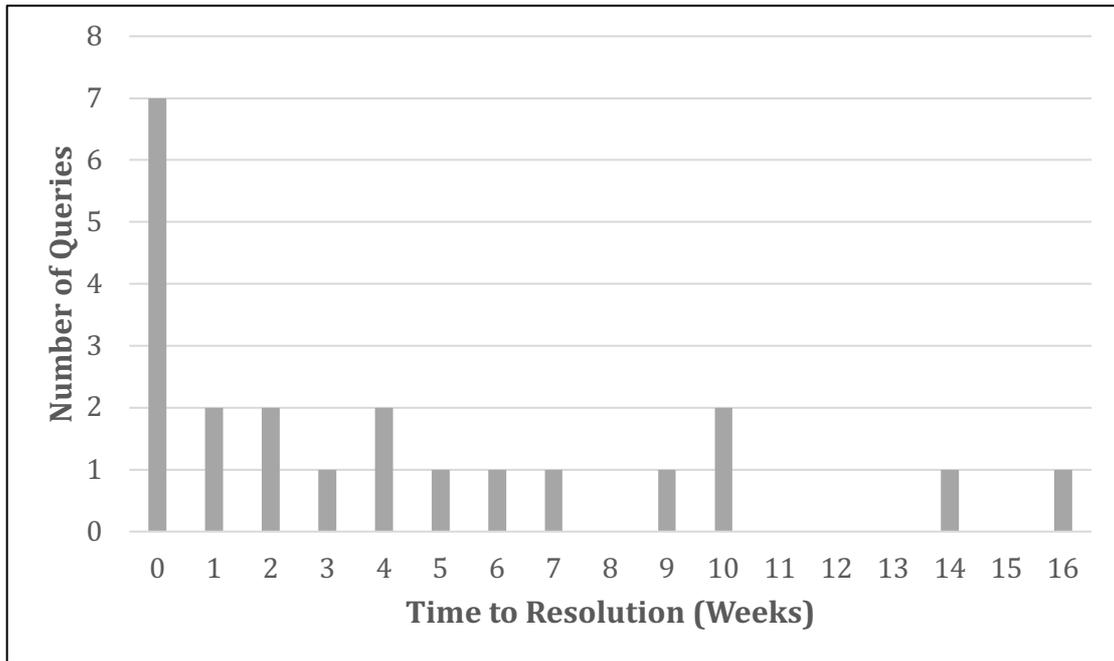
Q1216836 Help to compile and install the Qb@ll package on archer

A user needed to utilise the package Qb@ll, but this was not available on ARCHER. The users requested that the ARCHER CSE team install Qb@ll on ARCHER. Michael was able to install this package for them, resolving problems at the configuration stage and then with the integer size used by the Intel MKL interface. Michael was able to successfully test the code on ARCHER, before passing this back to the user.

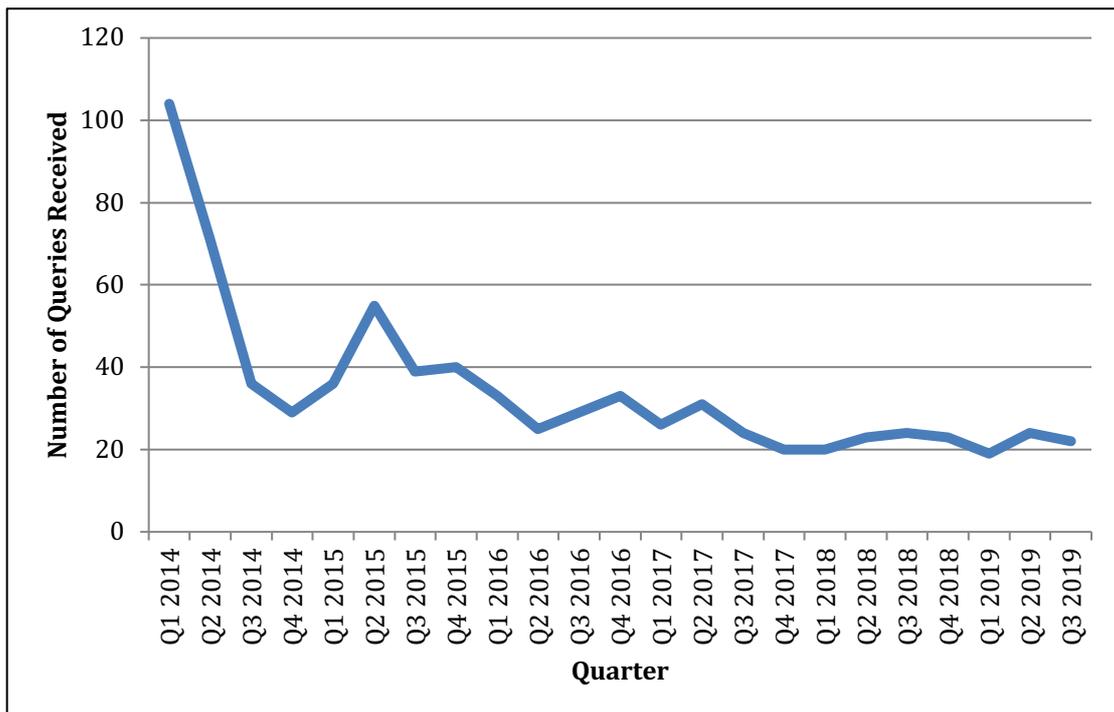
“To be honest, I have thought that, it might be not suitable to ask for this kind of support for third-party software from the archer team. However, I got a very nice support from you! Many thanks and all the best!”

In-Depth Query Analysis

The histogram below shows the time to resolution for In-Depth queries in the current reporting period.

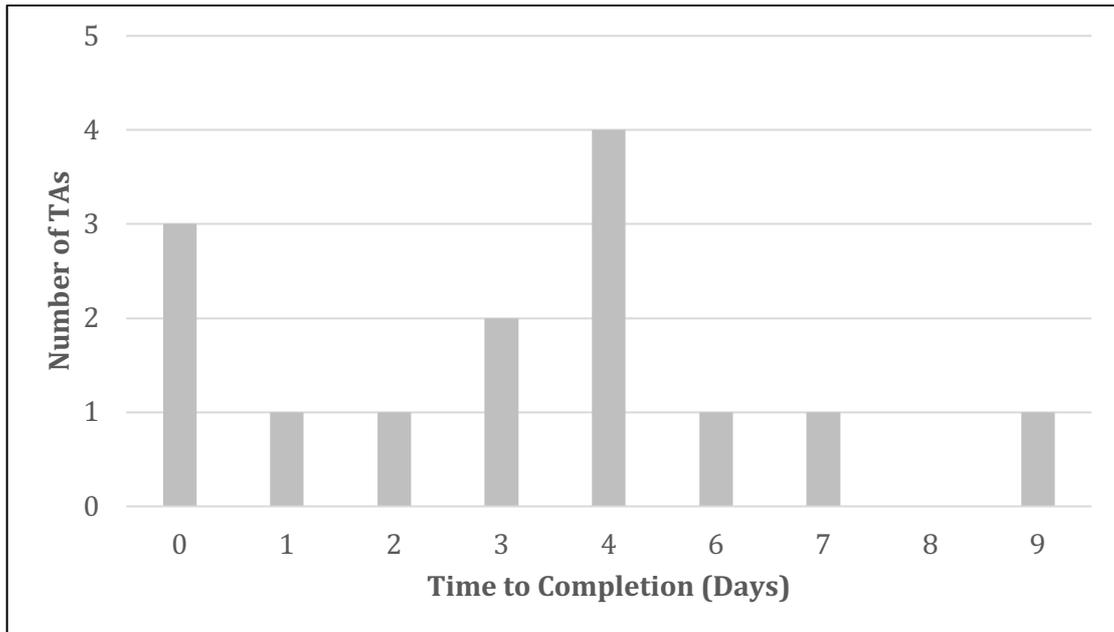


Plot of numbers of In-Depth queries received per quarter:

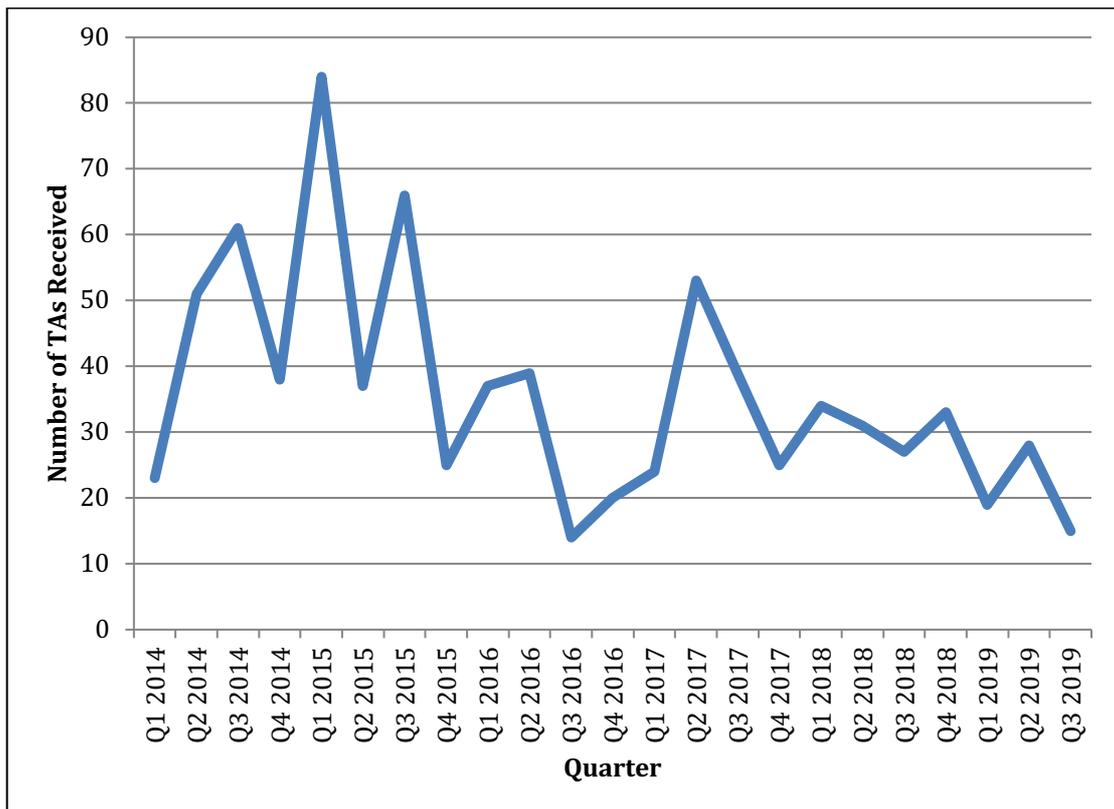


Technical Assessment Analysis

A histogram of the time to completion for Technical Assessments is shown below.



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.

Containers for HPC

Following on from our previous report, the investigation on using containers for HPC, with the focus on using Singularity containers (<https://www.sylabs.io/>), has been further progressed. We have developed a set of scripts to automatically provision a container suitable for the DiRAC-3 benchmarks. All supporting libraries and compilers have been successfully installed to a Singularity container. We are now testing the compilation of the benchmark applications. We plan to use ARCHER benchmarks for further investigations once we complete the current workflow investigation.

ARCHER Data Transfer Guide update

Also following on from our previous report, we have now completed the update of the ARCHER Data Transfer Guide and the new release has been available to ARCHER users. The updated ARCHER Data Transfer Guide provides best practice in transferring data between ARCHER, the RDF and remote machines, aiming to help users with data transfer toward the end of the ARCHER service. A virtual tutorial on “Good Practice for transferring data” was delivered on 4 September. The blog post, slides and video are all available on the website.

OpenFOAM new I/O model

As one of the most widely used applications on ARCHER, OpenFOAM is historically considered to scale poorly on large-scale HPC facilities. In particular, the way that the old versions of OpenFOAM handled I/O generated large numbers of files and caused huge amounts of I/O cost. This may also lead to the max number of files allowed per user on ARCHER being exceeded. The recent version OpenFOAM-5 introduced a new I/O model with collated file format and reduces the number of file sets from $O(N^2/2)$ to $O(N)$. We are now benchmarking this collated file format on ARCHER and also testing its efficiency against the old I/O method. We plan to make these benchmarks available to users and will provide guidance to users on the new I/O model and on the optimal settings for ARCHER.

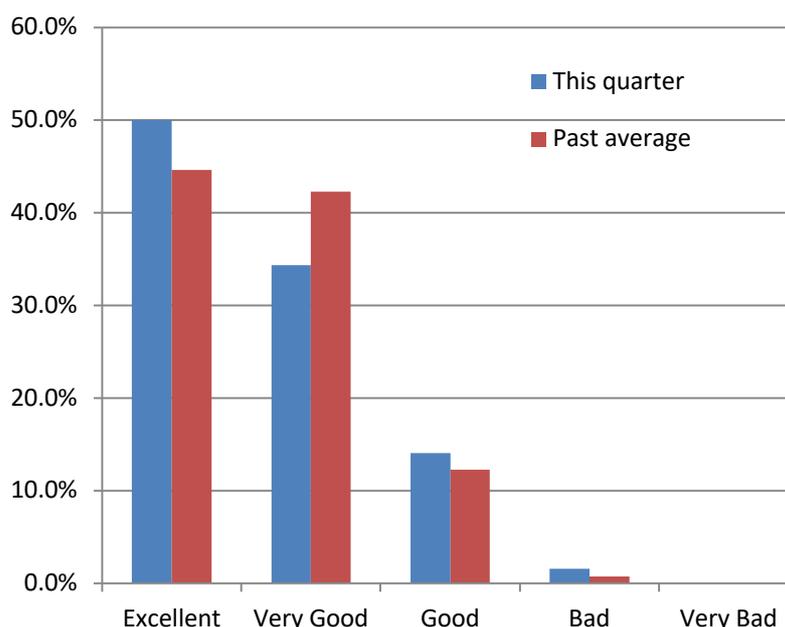
8. Training

This quarter, the CSE Service has provided a total of 21 days (378 student-days) of face-to-face training across 8 different locations and 1.5 days of online tutorials (average attendance 15 per tutorial).

Month	Dates	Course	Location	Days	Attend
Jul 2019	4-5	GPU Programming with CUDA	London	2	9
	10-11	Practical Software Development	Leeds	2	18
	15-16	Hands-on Introduction to HPC	Edinburgh	2	25
	17-19	Message-Passing Programming with MPI	Edinburgh	3	27
	22-23	Advanced MPI	Oxford	2	14
	22-26	2 nd Advanced VAMPIRE Workshop	York	1*	23
	30-31	Single Node Performance Optimisation	London	2	15
Aug 2019	7-9	Parallel Design Patterns	Durham	3	8
	27-28	Modern C++ for Computational Scientists	Glasgow	2	30
Sep 2019	28	OpenMP on GPUs	Online	0.5	
	4	Data Transfer from ARCHER	Online	0.5	
	25	Enabling distributed kinetic Monte Carlo simulations	Online	0.5	
	Sep 30 - Oct 1	Programming the ARM64 Processor #	Cambridge	2	14

* As part of this week-long consortium workshop, the CSE team provided one day of hands-on training on how to run the VAMPIRE software on ARCHER.

This formed part of our long-running series of Porting and Optimisation Workshops.



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 64 feedback forms, a response rate of 40%.

The one response of “Bad” was for the C++ course. The user did not leave their contact details but, looking at the feedback, the issue was that they felt the exercises were too difficult. We will review the format of the exercises, for example to consider whether they could be split up into smaller chunks with separate solutions supplied for each sub-exercise. We will also review the pre-requisites to ensure that attendees are fully aware of any assumptions about pre-existing programming experience. Overall, however, the C++ course had an average rating of 3.8 (just below “Very Good”) so there are no major concerns.

12 days of face-to-face training are planned for the fourth quarter of 2019, plus 3.5 days online. This will exceed our target of 72 training days for the year.

Month	Dates	Course	Location	Days
Oct 2019	16 & 23	LAMMPS Workshop	Online	1*
	29-30	Data Carpentry	Cardiff	2
	30 Oct – 1 Nov	Hands-on Introduction to HPC for Life Scientists	Birmingham	3
Nov 2019	31 Oct – 1 Nov	Introduction to Spark for Data Scientists	London	2
	13	Fully Lagrangian Dynamical Core for UKMO/NERC Cloud Model	Online	0.5
	13 Oct – 4 Dec	Shared-Memory Programming with OpenMP	Online	2*
Dec 2019	TBD	Containers for HPC	TBD	1
	9-10	HPC Carpentry	Edinburgh	2
	16-17	Scientific Programming with Python	Belfast	2

* These online interactive courses will be run in an extended format over consecutive Wednesday afternoons, giving time for attendees to attempt the practical exercises between sessions and to raise any issues with the trainer at the following webinar.

9. Embedded CSE (eCSE)

eCSE Scientific Highlights

As we approach the end of the ARCHER service, we continue to create and publish scientific highlights for completed projects. A small number of these have been selected to illustrate the eCSE work over the reporting period.

A fast coupling interface for a high-order acoustic perturbation equations solver with finite volume CFD codes to predict complex jet noise

Dr H. Xia, Prof G. Page, Dr M. Angelino and Mr M. Moratilla-Vega from Loughborough University

This eCSE project has developed an interface to help simulations looking to reduce aircraft noise. Aircraft noise is a major concern with strong negative health and economic impacts. This project has developed a fast interface for a high-order acoustic perturbation equations solver which couples with finite volume Navier-Stokes equations codes. This targets a computational issue that hinders more efficient modelling of sound waves leaving the turbulent jet stream. The resulting interface allows massive amounts of data to be exchanged in a way that is as parallel as possible. The resulting simulations are 10 times faster for small problems and 100 times faster for larger problems. Predicting the propagated far-field noise emitted from turbulent jets can now be done much more efficiently.

A fully Lagrangian dynamical core for the Met Office NERC Cloud Model

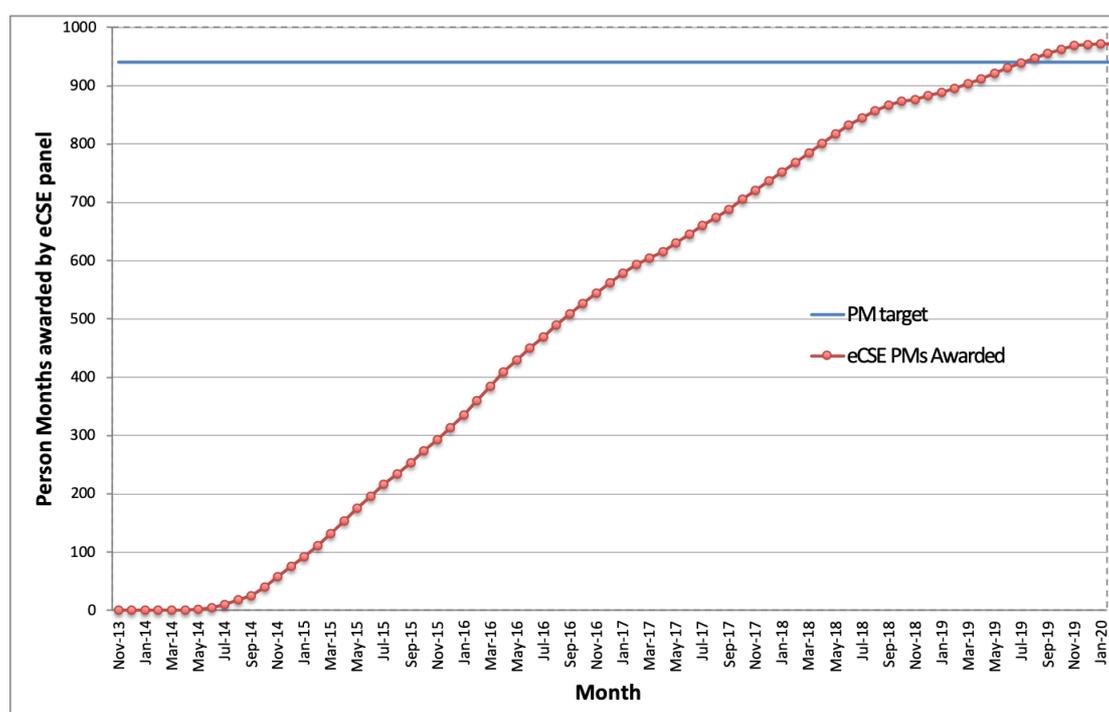
David Dritschel (University of St Andrews), Nick Brown, Michèle Weiland and Gordon Gibb (EPCC), Doug Parker, Alan Blyth and Steven Böing (University of Leeds)

The turbulent behaviour of clouds is responsible for many of the uncertainties in weather and climate prediction. Weather and climate models fail to resolve the details of the interactions between clouds and their environment and suffer from a crude representation of microphysical processes, such as rain and snow formation. To overcome these difficulties, the group have recently developed a prototype new method called "Moist Parcel-In-Cell" (MPIC) which deals with the dynamics of clouds in an essentially Lagrangian framework.

The overall aim of the eCSE project was to parallelise the Lagrangian 'Moist Parcel in Cell' (MPIC) code by incorporating it into the 'Met Office NERC Cloud Model' (MONC) framework. MPIC was parallelised only by OpenMP and as such could not scale beyond one node. MONC has MPI parallelisation and is known to scale to 30,000 cores.

This eCSE project has allowed the new code "the 'Parallel Moist Parcel in Cell", or PMPIC code, to scale to many thousands of MPI processes, permitting much larger simulations to be carried out than was originally possible with MPIC. Furthermore, modernisations to the codebase of MPIC achieved a single core speedup of 55%.

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 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	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	5	1 final report is late and being pursued.
eCSE11	18	6	49	6	6	4	1 late final report is being pursued. 1 late report will no longer be pursued.

eCSE12	23	6	41	6	6	5	1 late final report is being pursued.
eCSE13	21	10	96	10	2	0	1 late final report is being pursued and the remainder of the final reports are due by the end of the ARCHER Service.
Total	222	100	973	100	92	85	

- 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 has been reviewed by the eCSE Panel and appears to have been considered a success. This risk will now be removed from the risk list.
 - 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.
 - eCSE12-20: the project runs right up until 31 October 2018 – almost the end of the original CSE contract
 - The final report has been reviewed by the eCSE Panel and appears to have been considered a success. This risk will now be removed from the risk list.