



# ARCHER CSE Service Quarterly Report

Quarter 3 2017



## 1. Executive Summary

This report covers the period: 1 July 2017 to 30 September 2017 inclusive.

- Centralised CSE Team:
  - Comparative benchmarking across ARCHER and Tier-2 systems has progressed with the aim of helping users choose the best HPC service for their research. Results for CASTEP and HPC Challenge are now publicly available on GitHub: <https://git.io/vdr2Z>  
<https://git.io/vdr2C>
  - We have used new SAFE functionality to analyse the memory used by all jobs running on ARCHER to help with design of ARCHER2 . Initial results are interesting with most jobs having memory requirements below 1 GiB/core.
  - We led the development of the HPC-UK website (<http://www.hpc-uk.ac.uk>) in collaboration with UK-RSE and the Tier-2 HPC centres to help UK researchers make the most of the HPC ecosystem currently available.
- Training:
  - We delivered 16 days (345 student-days) of face-to-face training in the quarter at 6 different locations, with an average feedback score better than “Very Good”.
  - To ensure we respond to user demand, we leave 20% of training unscheduled at the start of each year. This enabled us to deliver three courses this quarter based directly on user requests: Introductory MPI in Exeter and York, and a bespoke Advanced MPI course for the UKCTRF consortium in Southampton.
  - We are taking an active role in the efforts to develop an “HPC Carpentry” course as we believe this will have substantial benefits for the UK HPC user community.
- ARCHER Outreach Project:
  - During June and July, we ran the EPSRC supported “Best use of ARCHER” competition that awarded early career (PhD or Post-Doctoral) ARCHER users with £3000 for travel to build collaborations between the US and UK. This facilitates early career researchers to build and develop their international network. The Awards were presented at an evening reception in London on 28th September, held jointly with the SSI US/UK Competition winners. As well as producing valuable publicity material, this provided an opportunity for the winners to network with those in the ARCHER community and with Research Software Engineers from across the UK.
  - Diversity in HPC website (<http://www.hpc-diversity.ac.uk>): we now have a total of 44 ‘Faces of HPC’ showcased on the website including 20 historical biographies and 24 interviews. This is now developing into a strong body of biographies, showcasing diversity and inclusion in HPC.
  - The Outreach team attended New Scientist Live, a large outreach event at the ExCel in London. During the three days, the booth had around 2000 visitors and provided an ideal opportunity to show the importance of HPC to the public.
- eCSE:
  - Of 84 projects from the first 11 eCSE calls, 82 have started and 63 have already completed; the remaining 2 projects are due to start within the next 6 months. Of those completed, 56 final reports have been received with 17 presently at various stages in the review process.
  - Proposals from the final call, eCSE12, are presently under review. We are expecting to award around 40 person months at this panel in addition to the extra person months already awarded.

## 2. Collaborations and Outputs Summary

- Presentations:
  - Andy Turner, Scaling the HPC Pyramid – RSE support for researchers on advanced computing facilities, Workshop, RSE 2017, 7-8 Sep 2017, Manchester, UK
- Meetings:
  - Neelofer Banglawala, Materials Chemistry Consortium (MCC) Annual Meeting, 5 Jul 2017, London, UK
  - Lorna Smith, Adrian Jackson, Andy Turner, Chris Johnson, ARCHER2 Town Hall Meeting, 12 Jul 2017, London, UK
  - Andy Turner, Manos Farsarakis, CSE/DiRAC/IPCC Workshop, 13 Jul 2017, Hartree Centre, UK
  - Andy Turner, RSE Leaders Group Meeting, 18 Jul 2017, Sheffield, UK
  - Andy Turner, DiRAC Day, 30 Aug 2017, Exeter, UK
  - Andy Turner, RSE 2017, 7-8 Sep 2017, Manchester, UK
  - Neelofer Banglawala, UK Consortium on Turbulent Reacting Flows (UKCTRF) Annual Meeting, 7 Sep 2017, Southampton, UK
  - Andy Turner, HPC-SIG Meeting, 13 Sep 2017, Hull, UK
  - Neelofer Banglawala, 5th CCP-BioSim Annual Conference, 13-15 Sep 2017, Southampton, UK
- Papers:
  - *A comparison of techniques for solving the Poisson equation in CFD*, Nick Brown  
In the journal of Civil Aircraft Design and Research (volume 3), pages 85-94

### 3. Forward Look

- Business Continuity and Disaster Recovery (BCDR): we are planning a live test of a simulated scenario for early Q4 2017:
  - This will provide important feedback on the performance of the current BCDR plan.
  - Lessons learned report will discuss ways to improve the current BCDR based on actual experience of putting it into action.
  - This is intended to be the first of an annual series of BCDR simulated tests.
- Centralised CSE Team:
  - Extend memory analysis to look at use by application and research area, and publish as an ARCHER white paper.
  - Update benchmarking exercise to cover additional applications and synthetic benchmarks, including the benchio parallel I/O benchmark.
  - Along with SSI and UK-RSE, we have had a BoF session titled “Software Engineers: Careers in Research” accepted for the SC17 programme.
- Training:
  - The ARCHER training plan for 2018 will be circulated at the start of October, to be considered by the Training Panel at a telcon on October 13<sup>th</sup>.
  - We will investigate new technologies for online training by attending the XSEDE online MPI course on 3-4 October as a trial satellite site.
  - We have had several presentations accepted for the SC17 programme including a BoF session “From Outreach to Education to Researcher - Innovative ways of expanding the HPC community”; a BoF session on “HPC Carpentry - Practical, Hands-On HPC Training” and presentations at the Workshop on Best Practices for HPC Training: “Creating Effective Learner Engagement in HPC Training and Beyond” and “HPC Carpentry”.
- Outreach:
  - Following the successful “Best Use of ARCHER” competition and celebration meal, we are in the process of editing short videos for a targeted section of the website. Together with the case studies, the aim is to showcase early career science on ARCHER and promote the service to other early career researchers.
  - WHPC has several workshops and activities at SC17; the theme is careers in HPC with an aim of directly improving the career opportunities of women in the community.
  - Following a successful and busy period of events, we plan to review feedback and input to enhance and improve demos

## 4. Contractual Performance Report

This is the contractual performance report for the ARCHER CSE Service for the Reporting Periods: July 2017, August 2017 and September 2017.

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	Jul-17		Aug-17		Sep-17		Q3 2017	
	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%	-4
<b>Total</b>		-5		-7		-7		-19

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

## 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	Jul-17		Aug-17		Sep-17		Q3 2017	
	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	Jul-17	Aug-17	Sep-17
Total Service Points	-6	-8	-8

## 5. 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 or enquiries about registration
<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 283 queries were resolved by the CSE service in the reporting period.

Metric	Jul-17	Aug-17	Sep-17	Total	% Total
Course Registration	80	61	52	193	68.2%
In-Depth	7	14	3	24	8.5%
Technical Assessment: RAP	0	1	13	14	4.9%
eCSE Application	2	2	10	14	4.9%
Technical Assessment: Grant	5	5	4	14	4.9%
Course Enquiry	10	2	1	13	4.6%
Technical Assessment: Instant	1	2	3	6	2.1%
Technical Assessment: HEC	0	4	1	5	1.8%

3 query feedback responses were received on In-depth queries in the reporting period. This represents a 13% return rate for feedback forms. All 3 responses registered a score of "Excellent". Work continues to implement SAFE functionality to send additional reminders to users on providing feedback responses to help increase the response rate.

Resolved In-Depth queries fell into the following categories:

Category	Number of Queries	% Queries
3rd party software	18	75.0%
Compilers and system software	3	12.5%
Login, passwords and ssh	1	4.2%
User programs	1	4.2%
Performance and scaling	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.

#### Q845505: internet access for jobs running on Archer

A user wanted to use a software application that runs remotely and collects data from GROMACS simulations running on ARCHER. This involves passing an external connection via SSH through the ARCHER login nodes and the PBS MOM nodes to the compute nodes running the GROMACS simulation. The CSE team provided a setup using the RSIP tools developed by the centralised team to allow this configuration to work. We expect to see more of these types of novel use of ARCHER for real-time data analysis in the future as more complex linking of multiple simulations occur.

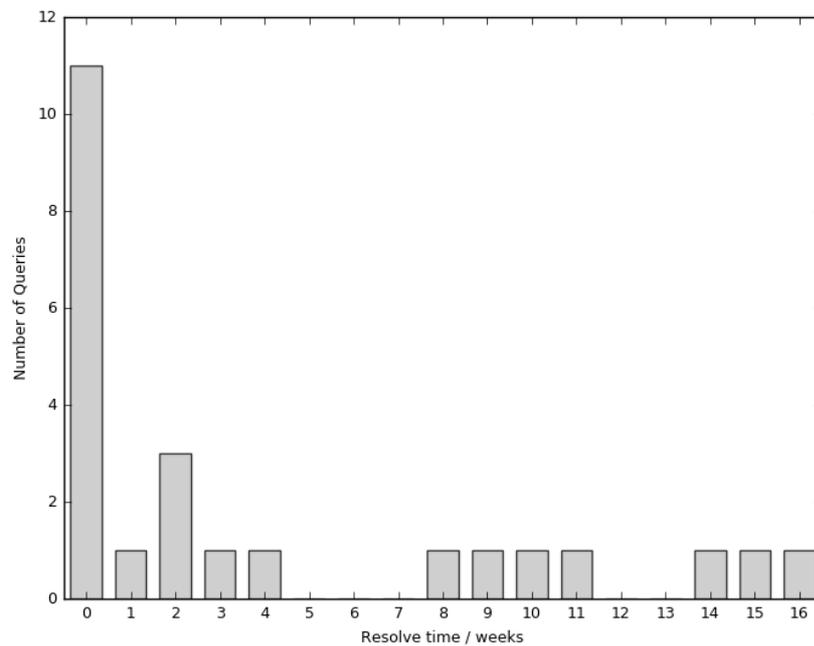
#### Q877041: TensorFlow on ARCHER

An ARCHER user wanted to use TensorFlow on the system and this required a number of dependencies that were complex to install and compile. The centralised team was able to untangle the range of dependencies and the complexity of the Python-based compilation system

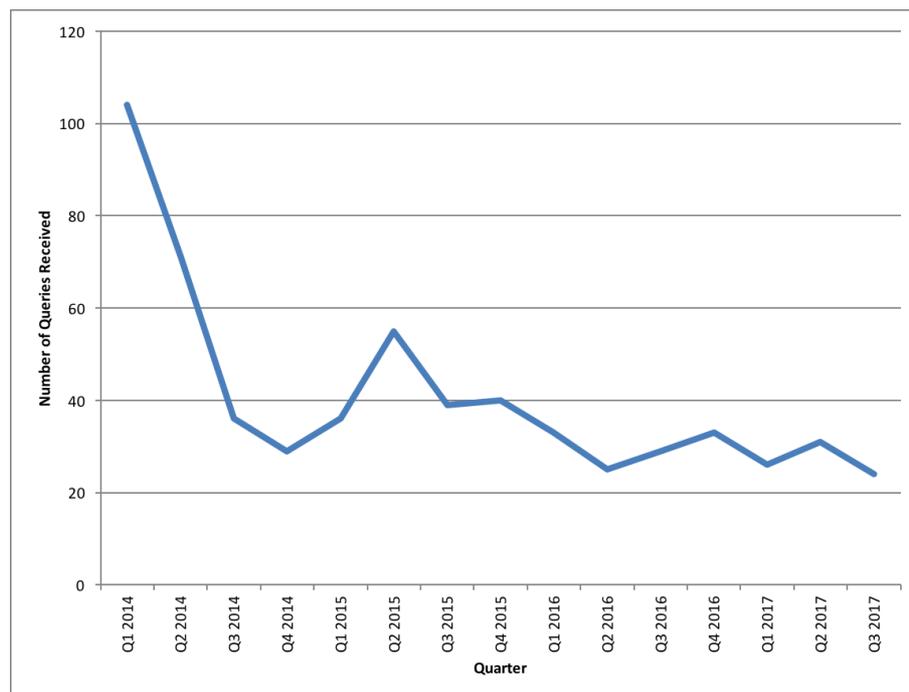
to create a parallel version of TensorFlow that works on top of the Cray MPI library. Having standard machine learning packages such as TensorFlow installed on ARCHER will allow data scientists to use the national supercomputer for their research problems.

### 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 weeks (median resolution time since 1 Jan 2014 is 2 weeks).

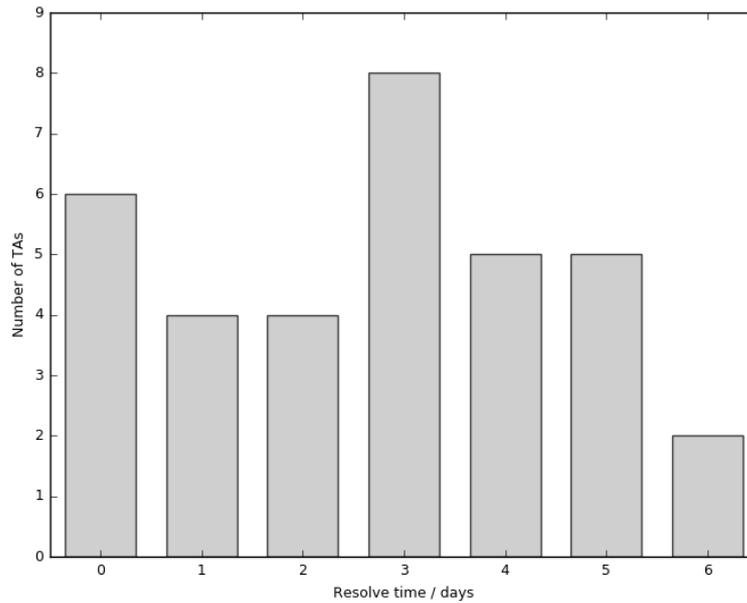


Plot of numbers of In-Depth queries received per quarter:

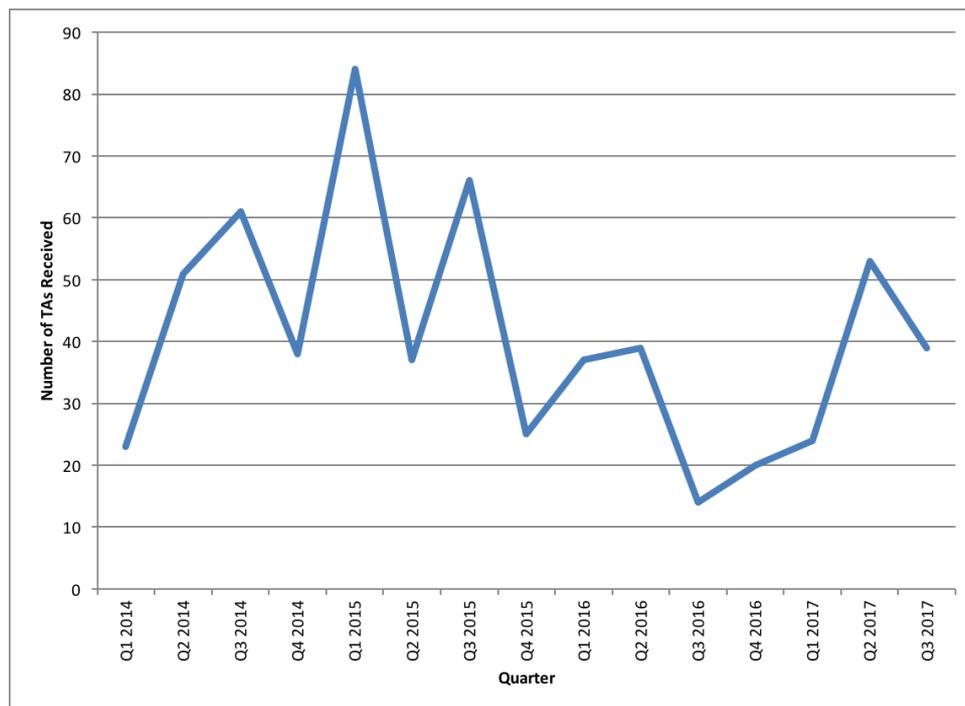


## Technical Assessment Analysis

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



Plot of numbers of Technical Assessments received per quarter:



## 6. Centralised CSE Team: Strategic Priorities Progress

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

### Future Technologies

#### ARCHER Benchmarks

We have extended the benchmarking process to the systems at HPC Midlands Plus (Athena) and MMM Hub (Thomas). So far, we have run a medium CASTEP benchmark and the bandwidth/latency test from HPC Challenge. We are currently unable to run the larger CASTEP benchmark on these systems due to restrictions in the job schedulers, and so have opened discussions with the systems to allow larger jobs to be run for benchmarking purposes.

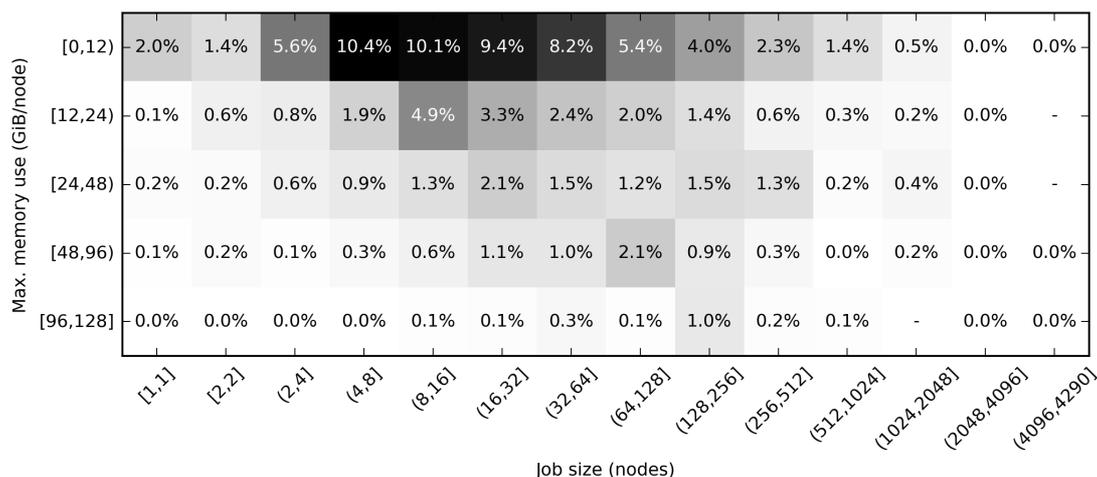
Initial analysis of the benchmark results can be found at:

- Medium CASTEP benchmark: <https://git.io/vdr2D>
- Large CASTEP benchmark (ARCHER and Cirrus only): <https://git.io/vdr29>
- Bandwidth/Latency: <https://git.io/vdr2Q>

We have recently gained access to the CSD3 system at Cambridge and will incorporate results from this system into our benchmarking results.

#### ARCHER Memory Use

We have updated the ARCHER SAFE to ingest data from the Cray RUR (Resource Usage Reporting) tool to allow us to measure additional information on how ARCHER resources are used by jobs. One of the metrics gathered by RUR is the maximum memory use across all processes within every job. We are using this metric to produce a report (with help from Simon McIntosh-Smith, University of Bristol) on the maximum memory use by jobs in the 1 year period from 1 July 2016 to 30 June 2017. A heatmap showing the maximum memory used per node as a function of job size is shown below.



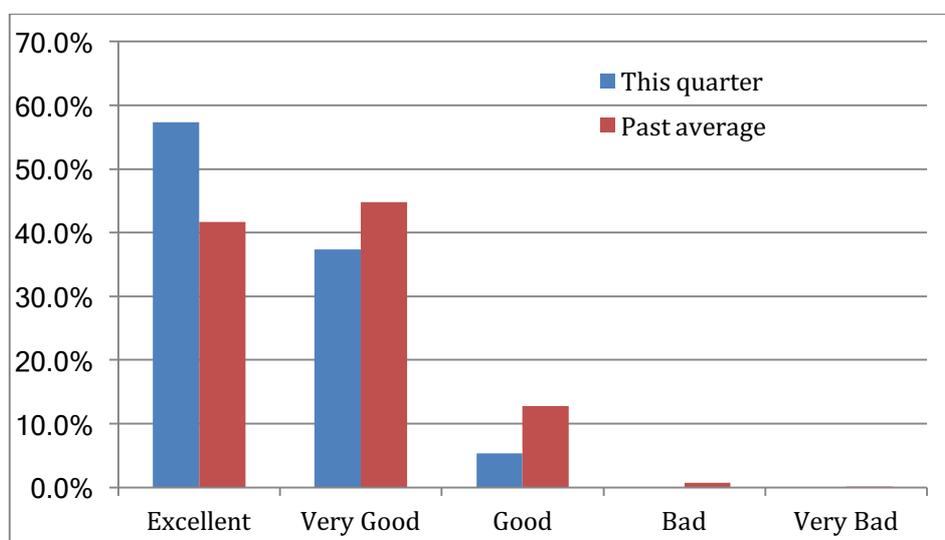
The data indicates that 60% of the usage on ARCHER uses less than 12 GiB/node (0.5 GiB/core) and 80% of the usage uses less than 24 GiB/node (1 GiB/core). This has potential implications for the design of ARCHER2. The full report will be published in Q4 2017.

## 7. Training

In the reporting period, the CSE Service has provided a total of 16 days (345 student-days) of face-to-face training across 6 different locations and 1.5 days of interactive online tutorials (average attendance 12 per tutorial).

Month	Dates	Course	Location	Days	Attendees
Jul 2017	10-11	Hands-on Introduction to HPC	Edinburgh	2	23
	12-14	Message-Passing Programming with MPI	Edinburgh	3	28
Aug 2017	19	Version Control	Online	0.5	
	27-28	Introduction to Modern Fortran	Cambridge	2	18
	1-2	Shared-Memory Programming with OpenMP	Oxford	2	11
	9	Interactive MPI Quiz	Online	0.5	
31 Aug - 1 Sep		Message-Passing Programming with MPI	Exeter	2	18
	Sep 2017	6	Advanced MPI	Southampton	1
11-12		Message-Passing Programming with MPI	York	2	28
12-13		Advanced MPI	Cambridge	2	30
	13	OpenMP 4	Online	0.5	

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



15 days of face-to-face training are planned for the final quarter of 2017, plus 1 day online.

Month	Dates	Course	Location	Days	Attendees
Oct 2017	11	OpenFOAM	Online	0.5	
	16-17	Software Carpentry	Glasgow	2	
	31 Oct -	Programming the Manycore	Cambridge	2	
Nov 2017	1 Nov	Knights Landing Processor			
	6-7	Single-Node Performance Optimisation	Oxford	2	
Dec 2017	8	Benchmarks of New HPC Hardware	Online	0.5	
	TBC	GPU Programming with CUDA	Daresbury	2	
	TBC	Data Carpentry	London	2	
	4-5	Hands-on Introduction to HPC	London	2	
	12-14	Advanced OpenMP	London	3	

## 8. Outreach Project

### ARCHER Best-Use Travel Competition

During June and July, we ran the EPSRC-funded “Best use of ARCHER” competition that awarded early career ARCHER users with £3000 for travel to build collaborations between the US and UK.

The winning entries covered a broad range of topics, including:

- helping to prevent pipeline blockages in the oil and gas industry;
- improving the performance of solar [power] by studying photovoltaic panel materials;
- and simulating combustion engines in order to improve efficiency and to reduce environmental impact.

The awards will facilitate these early career researchers, who are all either Ph.D. candidates or Postdoctoral researchers, to build and develop their international network. The competition aimed to identify the best scientific use of ARCHER, the UK’s national supercomputing facility, within the arena of the engineering and physical sciences. The Awards were presented at an evening reception at the Royal Geographical Society in London on 28th September in a ‘Celebration of new and enhanced international research opportunities in Computational Science’ providing an opportunity for the winners to network with those in the ARCHER community and with Research Software Engineers from across the UK. The winners will come together once again in 2018 to share with the supercomputing community the impact of the awards they received.

During the dinner, each of the winners had a short interview recorded. We are currently in the process of editing these to produce a series of short videos. Coupled with a case study, these will be showcased on the ARCHER web site.

### Diversity

Diversity in HPC website (<http://www.hpc-diversity.ac.uk>): we now have a total of 44 “Faces of HPC” showcased on the website including 20 historical biographies and 24 interviews.

Women in HPC:

- Skills to Thrive: Careers in HPC workshop held in collaboration with EuroMPI/USA 2017 at Argonne National Laboratories, Illinois, USA on 25 September 2017. This half day workshop, attended by 30 people and led by Women in HPC volunteers in the US, brought the WHPC “Skills to Thrive” series to the EuroMPI and Argonne National Lab communities. Topics covered in the workshop included perspectives on work-life balance, Evidence Based Interventions to Address Implicit Bias and Improve Workplace Climate, Discussing Approaches for Promoting Diversity and Inclusion at Workplace and finished with an open discussion on career development, mentoring and work-life balance.
- WHPC at the SC17 conference, November 2017: We are now planning our SC17 Women in HPC events. Throughout SC17 week WHPC will be focusing on a “Careers in HPC” theme, with the aim to directly improve the career opportunities of women in the community. Our events include:
  - Women in HPC@SC17 workshop: diversifying the workforce: a full-day workshop with two themes: providing women with the skills to thrive, and helping employers improve the workplace to attract and retain a diverse workforce.
  - Women in HPC Mentorship scheme: as part of SC17 our early career workshop presenters are all being offered the opportunity to be provided with a mentor to prepare for their presentation and making the most of SC. Mentors are all being offered training by WHPC.
  - Two BoF sessions on: “Non-Traditional Paths to HPC and How They Can and Do Enrich the Field” and “Recruitment: how to build diverse teams”. These

workshops will reach different groups in the community to help improve workplace inclusivity.

- Women in HPC evening reception: in collaboration with sponsors, WHPC will run an evening reception during the week of SC17. The evening reception will provide crucial networking opportunities for women in the community with a careers theme to bring together potential employers with women.

## Outreach Activities

It has been a busy time for the outreach team with five events across the UK and internationally. An estimated 2000-3000 visitors were seen by the team and gained insight into the relevance and importance of Supercomputing. Two key events are highlighted below:

Event	Location	Dates	Number of Visitors	Notes
New Scientist Live 2017	ExCel London	28th Sep - 1st Oct	1200-2500	Open to the public, this was a busy event and provided a good opportunity to explain the importance of Supercomputing to the public.
COMAC innovation event	Shanghai, PRC	4th - 8th September	150	Wee Archie travelled to Shanghai (with Nick Brown) to take part in an innovation event with others from UoE. We had a stand, with hand-outs & Wee Archie. The event showcased ARCHER to the international community and allowed sharing of best practice.

Ambassador pack: The contents have been modified based on team feedback. Most activities are now ready for publication.

## Impact

A new case study has been developed with the around protein shapes. This is in the final stages, just awaiting final approval from the researchers.

10 'mini' case studies have been completed for each of the winners of the US-UK competition and a booklet produced – several of these are also good candidates for full case study coverage.

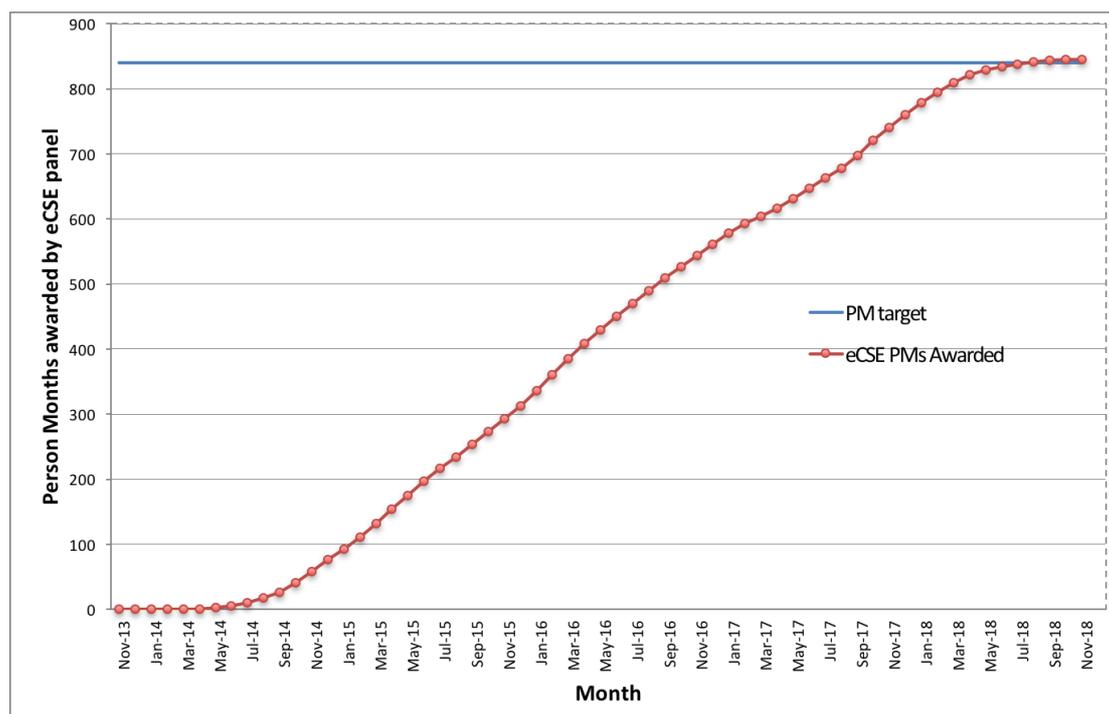
A review has been undertaken of the currently existing case studies and their coverage of EPSRC/NERC science areas. In general, there was good coverage of most areas, and we are working to identify any good candidates to cover any missing areas.

## Engagement

We are currently in the process of finalising the location for the next ARCHER Champions meeting, which is scheduled for Easter 2018.

## 9. Embedded CSE (eCSE)

### Overview of eCSE Effort



- The eCSE person months awarded up to and including the 11<sup>th</sup> eCSE call are shown in red.
- We committed to awarding at least 840 person months by the end of the project (14 FTEs for 5 years).
- 846 person months have been awarded so far over 84 awarded eCSE projects, with an expectation of around an extra 40 person months to be awarded at the final call which is presently at the reviewing stage.

### eCSE Call 1 – Call 12

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	9	1 late final report is being pursued.
eCSE04	16	8	82	8	8	7	1 late final report is being pursued.
eCSE05	14	8	94	8	8	8	
eCSE06	9	5	47	5	4	4	
eCSE07	16	5	49	5	5	3	2 projects have recently finished but final reports not yet due
eCSE08	21	8	88	8	5	2	2 late final

							report are being pursued. 1 other project has recently finished but the final report is not yet due
eCSE09	19	5	62	5	0	0	
eCSE10	13	6	65	6	0	0	
eCSE11	18	6	49	4	0	0	
eCSE12	23	N/A	N/A	N/A	N/A	N/A	Proposals under review
<b>Total</b>	<b>201</b>	<b>84</b>	<b>846</b>	<b>82</b>	<b>63</b>	<b>56</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:
  - eCSE04-4: the person named to do the technical work was offered a position elsewhere
    - The member of staff originally named on the contract completed 1.5 of the 12 months of work before leaving to take up another post. With approval from the PI and eCSE Panel chair, we identified a new member of staff within the ARCHER CSE team who took on the work from 01/10/15. This project is now complete and the final report has been received. The report has received a favourable technical review and is awaiting panel review.
  - eCSE04-10: the PI indicated that the person named to do the technical work may not be available
    - This project will go ahead with the original staffing. There was a short delay to the start of the project which started on 01/01/16. The project is now complete and we are awaiting the final report which is now overdue.
  - eCSE08-9: this project had a change of staffing
    - the new staff member has been approved by the panel chair. The project has started and the contract was signed. We will continue to monitor the work to the end of the project which is due to finish within the next reporting period.
  - eCSE08-10: there were issues raised by Cambridge University involving the IP and the relationship with the CASTEP group
    - these issues appear to have been resolved and a contract has been signed but we will monitor this as the project progresses.
  - eCSE09-6: this project will terminate 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 will resign from his current post in Cambridge sometime in November
    - the PI confirms that the first work package is likely be completed. Given the circumstances we have agreed to this early termination and any unused funds will be used to award eCSE12 projects.
  - 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. In addition, the contract is awaiting approval by the University of Reading.
    - of the 19 months awarded for this project, 7 are for a member of the ARCHER CSE team and the work will be monitored through EPCC's standard project monitoring processes. The remaining 12 are for an external member of staff at the PI's institution and will be monitored via regular contact with the PI.

- eCSE10-1 has been identified as being of medium risk due to the process of signing the contract having not yet been agreed
  - We will continue to pursue this via the legal teams within the University of Edinburgh and STFC
- eCSE10-5 was identified as being of medium risk due to a change of staffing being required
  - We have discussed this with the PI and it is likely that the project will be scaled back and re-staffed (subject to approval of the Panel chair) but will monitor the situation via regular contact with the PI. Any unused funds will be used to fund eCSE12 projects.
- eCSE10-10 has been identified as being of medium risk due to the process of signing the contract having not yet been agreed
  - We will continue to pursue this via the legal teams within the University of Edinburgh and STFC

### Future eCSE Calls

- We intend to use the remaining eCSE funds to award projects from the eCSE12 call and no future calls are presently planned.