



ARCHER CSE Service Quarterly Report

Quarter 1 2018



1. Executive Summary

This report covers the period: 1 January 2018 to 31 March 2018 inclusive.

- Centralised CSE Team:
 - We have completed the initial comparison of performance of ARCHER and the EPSRC Tier2 HPC systems for a range of synthetic and application benchmarks. We are currently writing a white paper on this work and will publish it shortly. This information will help users choose the most appropriate HPC resource for their research.
 - The CSE team has worked with the other ARCHER service partners to evaluate the potential performance impact of the mitigation for the Meltdown bug on ARCHER applications to evaluate the implications for the service of applying the mitigation. Our investigations indicate the impact on HPC workloads should be very small.
 - We have worked with NCAS and Cray to implement and test the use of ARCHER compute nodes for Unified Model compilation tasks. This will lead to a more reliable service for this key part of the UM workflow.
- Training:
 - We delivered 10 days (209 student-days) of face-to-face training in the quarter at 5 different locations, with an average feedback score better than “Very Good”
 - For the first time we ran a 2-day MPI programming course fully online, over four consecutive Wednesday afternoons. This was intended to make HPC training as widely available as possible and was a great success with some 60 users registering.
 - To raise awareness of the Tier1 and Tier2 systems available in the UK, we ran a virtual tutorial on “Getting Access to UK Advanced Computational Facilities for your Research” which attracted 28 attendees.
 - In order to promote the benefits of HPC as widely as possible, we publicised the recent run of the free “Supercomputing” MOOC to all ARCHER users.
- ARCHER Outreach Project:
 - For the third year running, we had a booth at the Big Bang Fair at the NEC in Birmingham from 14-17 March. The booth was very popular, and we estimate that over the 4 days we interacted with around 10,000 participants. We used the booth to encourage young people to consider careers in computer science, while explaining the benefit of Supercomputing to the wider public.
- eCSE:
 - Of the 90 projects awarded over the 12 eCSE calls issued, 89 have started and 72 have now completed. The one remaining project is due to start at the beginning of the next quarter and all are planned to finish by end of October 2018. Of those completed, 58 final reports have been received, 55 of which have been reviewed.
 - Our original target was to award 840 person months but we have been able to award an extra 41 person months giving 881 person months awarded in total.
 - We have now awarded projects using the entire eCSE budget and no further eCSE calls are planned within the current CSE contract.

2. Collaborations and Outputs Summary

- Presentations:
 - Rupert Nash, *HPC vs Cloud Case Studies*, 6 Feb 2018, MSc in HPC guest lecture, University of Edinburgh
 - Neelofer Banglawala, *Supercomputing*, Edinburgh Women in Technology, 28 Feb 2018, University of Edinburgh
 - Rupert Nash, *ARCHER Update*, UK Consortium on Mesoscopic Engineering Meeting, 28 Mar 2018, UCL
 - Andy Turner, *UK HPC Benchmarking*, 26 Mar 2018, HPC@Sheffield, University of Sheffield
- Meetings:
 - Andy Turner, *UK RSE Committee Meeting*, 10 Jan 2018, Online
 - Andy Turner, *HPE Technical Discussion*, 16 Jan 2018, University of Edinburgh
 - Andy Turner, *International RSE Leaders Workshop*, 30-31 Jan 2018, Alan Turing Institute
 - Andy Turner, *Intel SGX Technical Meeting*, 7 Feb 2018, University of Edinburgh
 - Andy Turner, *HPC-SIG Meeting*, 13 Feb 2018, University of Bristol
 - Kevin Stratford, *Programming multi-core and many-core systems using Kokkos*, 21-23 Feb 2018, University of Cambridge
 - Andy Turner, *Public Cloud for HPC Technical Meeting*, 23 Feb 2018, University of Edinburgh
 - Andy Turner, *SESC Advisory Board Meeting*, 8 Mar 2018, Online
 - Andy Turner, *UK RSE Committee Meeting*, 14 Mar 2018, Online
 - Amy Krause, Rosa Filgueira, *RSE for Data Science 2018*, 19-20 Mar 2018, Alan Turing Institute

3. Forward Look

- Centralised CSE Team:
 - We will extend our benchmarking exercise to include Tier2 GPGPU systems (CSD3 and JADE) and to ARM systems when available (Isambard) to help users decide if these different architectures could potentially meet the HPC requirements for their research.
 - To try and improve the service for users with heavy I/O requirements, we are analysing the variation in performance seen on the ARCHER Lustre file systems and working to understand what gives rise to this variation.
 - We are planning to extend our cross-platform benchmarking work to include machine learning (ML) benchmarks. ML is becoming more and more popular as a tool in research and understanding the performance of different HPC systems will allow new communities to exploit the national HPC facilities effectively.
- Training:
 - The “Data Analytics with HPC” and “Hands-on Intro to HPC” courses in Q1, which had to be postponed due to industrial action and adverse weather respectively, are being rearranged for Q2. Users registered for the initial runs will be given priority when registering for the new dates.
 - We have collated input from the fifth “Training Impact Survey” and will circulate a document describing the results in April.
 - We are continuing our collaborations with the Alan Turing Institute (ATI) by running “Scientific Programming with Python” at the ATI in London. For ATI students, this course will form part of their Researcher Training Programme.
- Outreach:
 - The outreach grant is now finished. In future, we plan to focus on ensuring activities and material developed during the grant is exploited fully. For example, through promotion of the teachers pack and ambassadors pack.

4. Contractual Performance Report

This is the contractual performance report for the ARCHER CSE Service for the Reporting Periods: October 2017, November 2017 and December 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	Jan-18		Feb-18		Mar-18		Q1 2018	
	Perf.	SP	Perf.	SP	Perf.	SP	Perf.	Total
QE1	100%	-2	100%	-2	100%	-2	100%	-4
QE2	80%	0	100%	-2	100%	-2	100%	-6
TA1	100%	-1	100%	-1	94%	-1	100%	-3
FB1	100%	-2	100%	-2	100%	-2	100%	-6
Total		-5		-7		-7		-21

Pink – Below Service Threshold

Yellow – Below Operating Service Level

Green – At or above Operating Service Level

QE2 was below Operating Service Level in January 2018 due to a single In-Depth query taking longer than 4 months to resolve. As the number of queries in a period is low (5 queries in January 2018), a single query has a large impact on statistical measures.

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-18		Feb-18		Mar-18		Q1 2018	
	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-18	Feb-18	Mar-18
Total Service Points	-6	-8	-8

5. 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 430 queries were resolved by the CSE service in the reporting period.

Metric	Jan-18	Feb-18	Mar-18	Total
Course Registration	142	160	55	357
eCSE Application	0	4	1	5
In-Depth	5	6	9	20
Course Enquiry	4	4	6	14
Technical Assessment: Grant	5	6	5	16
Technical Assessment: Instant	1	4	1	6
Technical Assessment: RAP	0	0	12	12
Total	157	184	89	430

2 query feedback responses were received on In-depth queries in the reporting period. This represents a 10% return rate for feedback forms. 1 response registered a score of "Excellent" and 1 response registered a score of "Good". We continue to try to improve the response rate for feedback from queries by offering charity donations for responses and sending an additional reminder to users to provide feedback.

Resolved In-Depth queries fell into the following categories:

Category	Number of Queries	% Queries
3rd party software	13	75%
Compilers and system software	2	10%
Porting	2	10%
Other	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.

Q958952 Running on multiple nodes on conda version of MPICH

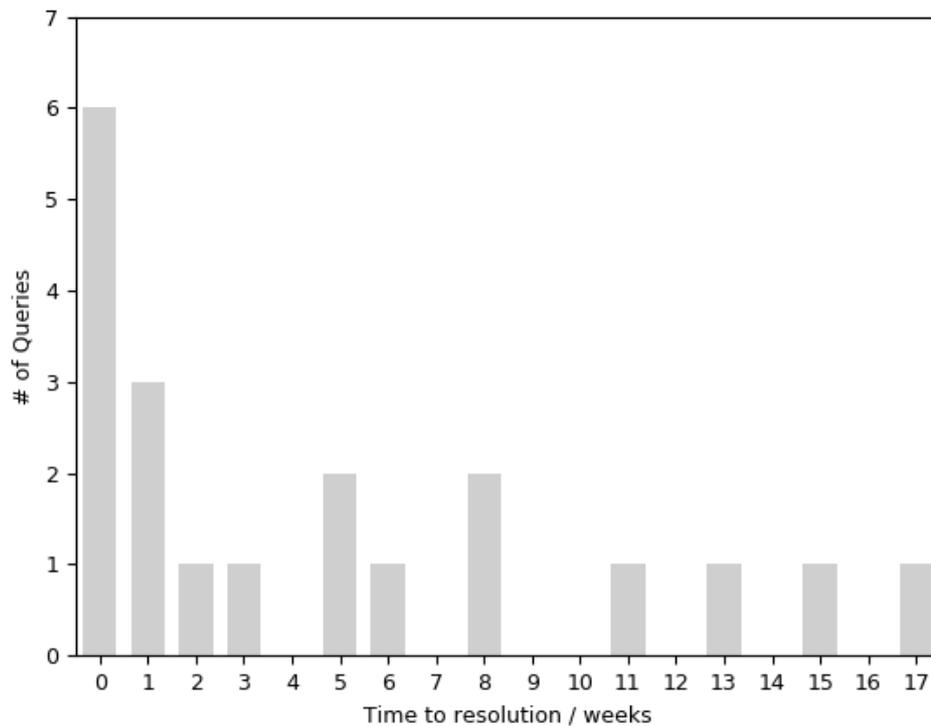
A RSE working to develop and deploy a coupled application on ARCHER (and other HPC systems) was trying to put together a simplified deploy system for their software. The compilation procedure is complex and a simplified deploy using the conda Python setup would ease the process for researchers to use the software on ARCHER. The conda deployment option provided its own version of MPI which was incompatible with the ARCHER interconnect leading to problems using more than one compute nodes for jobs. After the issue had been described to the RSE, the CSE team worked with them to gather their requirements for simplified deployment on ARCHER and provided a number of alternatives that would work with ARCHER along with their respective advantages and disadvantages. Using this expert advice, the RSE was able to produce a plan to take the deployment work forwards.

Q849753 extended use of rsip on a login node

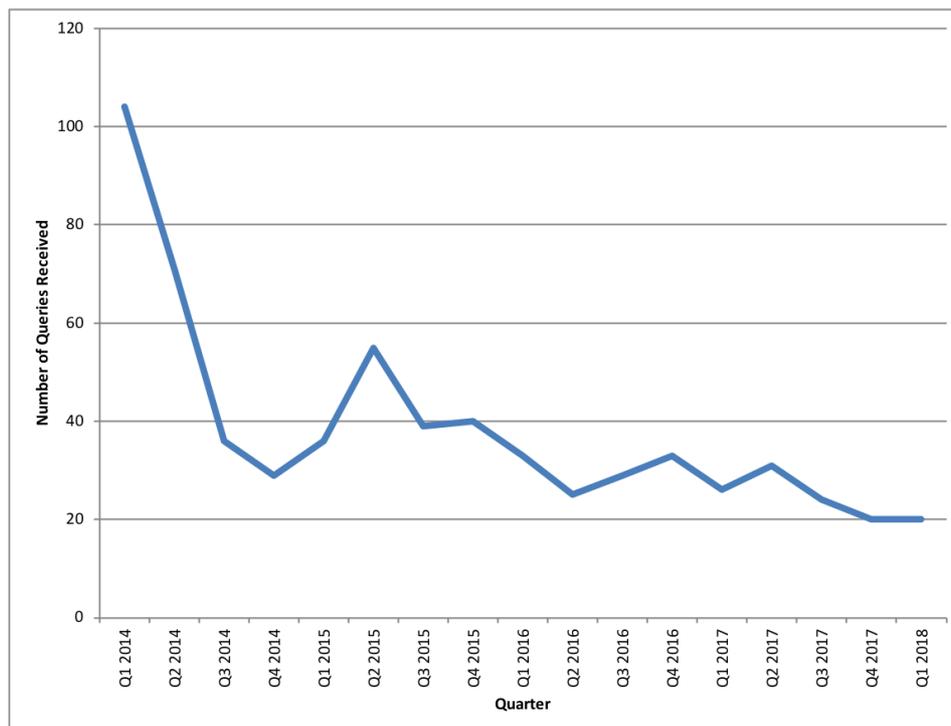
An ARCHER user wanted to run simulations on ARCHER that periodically communicate information on the simulation back to a database running at their local institution. As the ARCHER compute nodes cannot talk directly to the external internet, the user wanted to know how they could make this work. The CSE team had previously helped the user use the ARCHER login nodes to enable such traffic but this mechanism is difficult to keep alive for the duration of the job. We created a solution that used the serial, post-processing nodes as the RSIP bridge between the compute nodes and the external internet; this allowed the connection to be maintained for the duration of the job running on the compute nodes enabling the user to proceed.

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

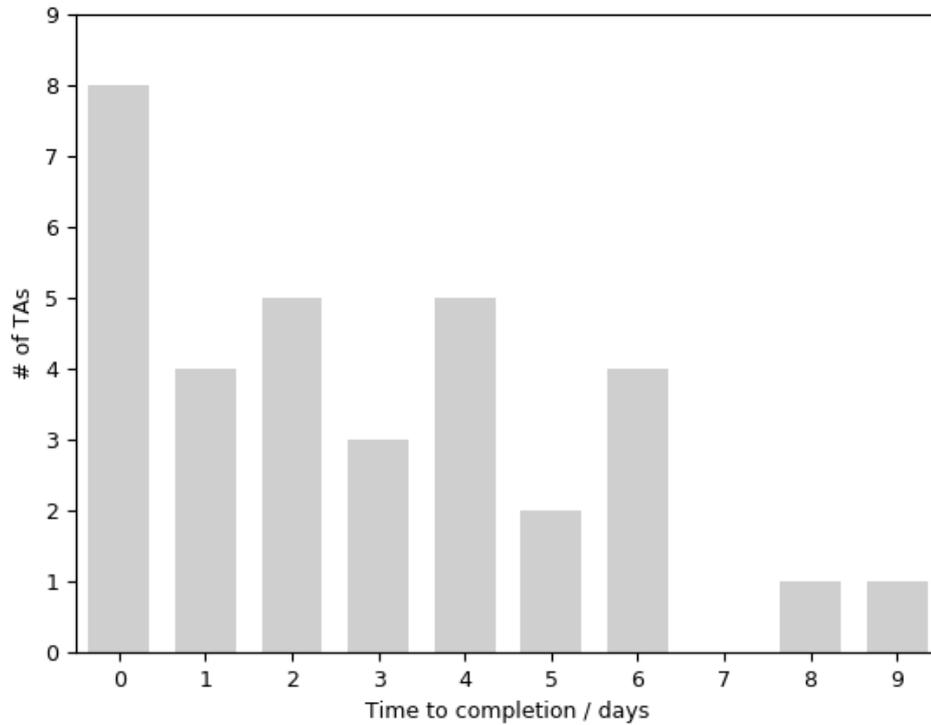


Plot of numbers of In-Depth queries received per quarter:

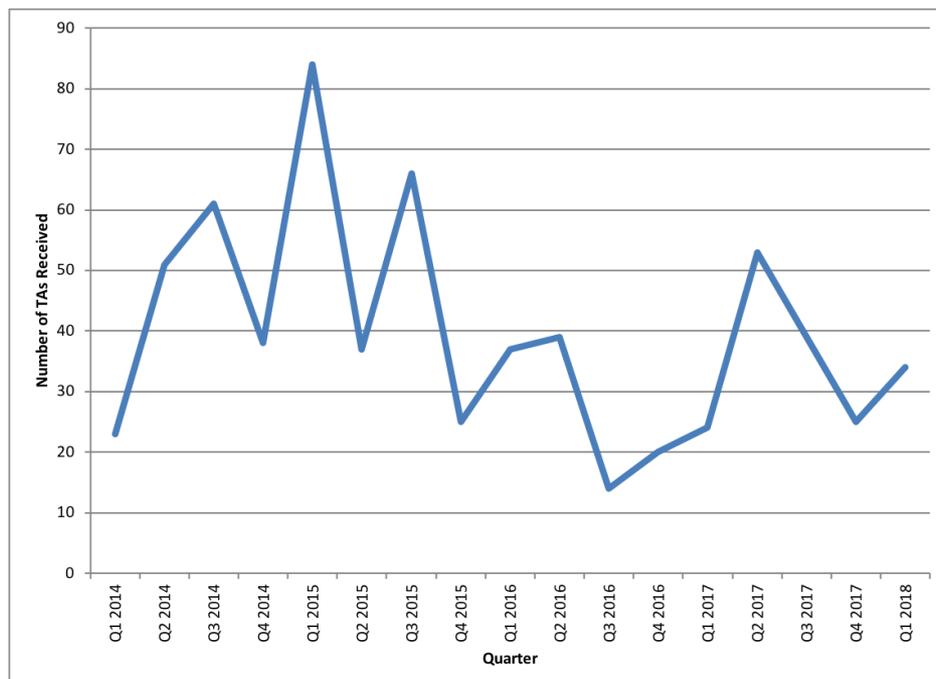


Technical Assessment Analysis

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



Plot of numbers of Technical Assessments received per quarter:



6. 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.

Comparing UK HPC Systems

We have compared the performance of ARCHER and a number of the EPSRC Tier2 HPC systems using the set of ARCHER application and synthetic benchmarks that were produced by the CSE service in collaboration with the ARCHER user community.

In this initial comparison, we included all the Intel Xeon based HPC systems (comparisons including GPGPU and ARM systems will be added in the future). The systems included three different generations of Xeon processor technology, multiple types of interconnect, and Lustre and IBM SpectrumScale parallel file systems.

The comparison reveals that the performance of compute- and memory bandwidth-bound application benchmarks generally increases for the newer generation of processors while applications that are bound on random memory access do not see significant improvements. Where we are able to scale out to large core counts, we find that the ARCHER interconnect generally outperforms those on the Tier2 systems. Parallel file system benchmarks show a large variation in performance between different systems.

We plan to extend these comparisons in the future by including the Tier2 GPGPU systems, and also plan to approach the DiRAC service for access to their systems to include in the comparison.

Investigating metadata server performance

We have previously focussed on the write bandwidth as a metric for assessing the performance of parallel file systems as this is the key metric for well-written applications performing large parallel I/O operations. However, for many use cases on parallel file systems the performance of the metadata servers (MDS) can be the limiting factor. Generally, the MDS performance is critical to operations such as creating/deleting files and directories and getting the status of files. If the MDS gets overloaded with such requests (from all users on the system) then the MDS response time worsens, often resulting in very slow response times to the operations described above.

We have recently begun to investigate the MDS performance on ARCHER and other HPC systems using the *mdtest* benchmark that was used in previous large US supercomputing procurements. The initial results show that the ARCHER MDS performance is significantly lower than that on other systems (e.g. Cirrus, CSD3); we are working with Cray to understand these results and plan to publish the results in an upcoming ARCHER white paper.

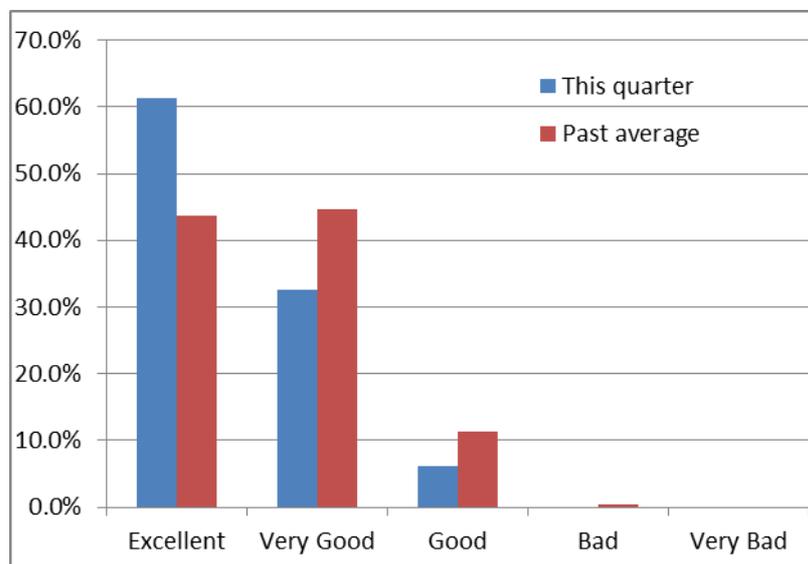
7. Training

The CSE Service has provided a total of 10 days (201 student-days) of face-to-face training across 5 different locations plus 2 days of interactive online tutorials (average attendance 33 per tutorial). There were also 2 days of the new Online MPI course - we estimate that around 50 people attended the first day, dropping to around 25 in the final week.

Month	Dates	Course	Location	Days	Attendees
Jan 2018	24	Parallel SPH (eCSE project)	Online	0.5	
	31	MPI online (on four consecutive Wednesday afternoons)	Online	2	32*
Feb 2018	8	Porting and Optimisation Wshop	Oxford	1	10
	26-27	Software Carpentry	Leicester	2	25
	27-28	OO Programming in Fortran	Daresbury	2	9
Mar 2018	28	State of the Art IO Tools	Online	0.5	
	7	Getting Access to UK Advanced Computational Facilities	Online	0.5	
	20-22	Threaded Programming	Soton	3	35
	20-21	Efficient Parallel IO	Cambridge	2	9
	28	Meltdown for Dummies	Online	0.5	

* Although we can easily monitor the number of attendees at online courses, several UK sites set up dedicated rooms where groups of users could attend, watch the presentations and do the practicals together. These only show up as a single connection so this attendance figure is an estimate, as opposed to face-to-face courses where figures come from a physical sign-in sheet.

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 49 feedback forms, a response rate of 56%.



19 days of face-to-face training are planned for the second quarter of 2018, plus 1.5 days online.

Month	Dates	Course	Location	Days	Attendees
Apr 2018	11-13	Message-Passing Programming with MPI	Soton	3	
	11	TBC	Online	0.5	
May 2018	26-27	Advanced MPI	Exeter	2	
	9	TBC	Online	0.5	
Jun 2018	10-11	Hands-on Intro to HPC	Edinburgh	2	
	13	TBC	Online	0.5	
	17-18	Scientific Programming with	London	2	
	TBC	CP2K Workshop	Daresbury	3	
	TBC	Software Carpentry	London	2	
	TBC	Data Analytics with HPC	Belfast	2	
	TBC	Performance Analysis Workshop	London	3	

8. Outreach Project

This is the last quarter of the ARCHER Outreach project as the grant finished at the end of March. In future, we will look to maximise the impact of the various outputs produced during the grant. For example, we will continue to utilise Wee Archie at various outreach events, continue to contribute to our leadership role in promoting diversity and inclusion in HPC, and encourage the uptake of both the Ambassadors and Teachers packs. We are running a further ARCHER Champions meeting in April.

Engagement

The fifth ARCHER Champion Workshop will be in Manchester on 25th April. It is being collocated with the SSI Workshop: Impact of International collaborations in research software. Topics will include the eCSE programme, training and MOOC, Tier2 as well as a lightning session with contributions from Champions. The future of Champions will also be discussed at the workshop.

Diversity

Having previously developed guidance on ensuring training material is accessible, all ARCHER training courses have now been updated to meet this guidance. This is an important part of ensuring our resources and services are available to all.

We have appointed Manos Farsarakis as our diversity coordinator, replacing Toni Collis. Toni has left EPCC but is still involved in Women in HPC and is working with us to ensure the on-going success of Women in HPC. Manos is involved in a submission to ISC 2018, in collaboration with Women in HPC, for a BoF on Diversity. He will also present at the next ARCHER Champions.

Outreach

For the third year running we had a booth at the Big Bang Fair at the NEC in Birmingham on the 14-17 March. Activities included the bean-bag sorting activity, the supercomputing app and Wee Archie. In addition, we had an actual XC30 blade, courtesy of Cray, and full-size photographs of an ARCHER XC30 cabinet. This included the inside of the cabinet from the front (showing the blades) and the back (showing the cabling). This allowed attendees to see what ARCHER actually looks like and gain a sense of the scale of the machine.

The booth was very popular, and we estimate that over the 4 days we interacted with around 10,000 participants. Attendees came from all over the UK with the first three days only open to school groups, with many of these children less familiar with science events.

Wee Archie was part of the EPSRC booth at the American Association for the Advancement of Science exhibition in Texas. This provided an opportunity to showcase the UK's world-leading HPC outreach activities that have been supported through the ARCHER Outreach grant.

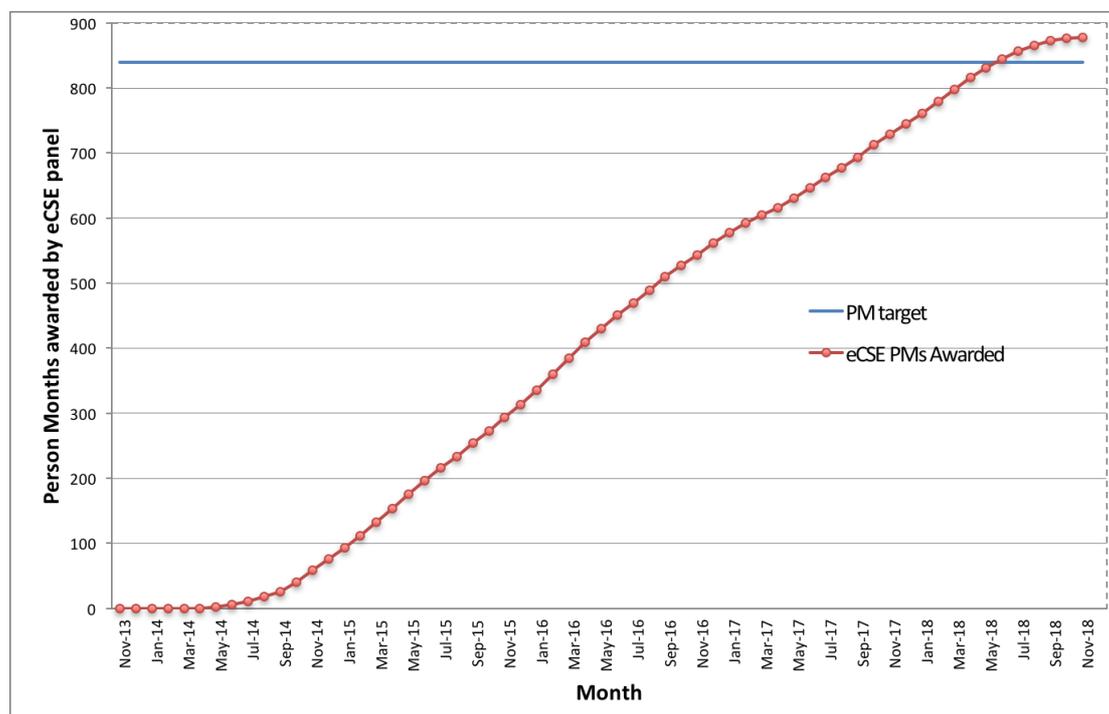
The Ambassador Pack is now complete and published. This resource is designed to support ARCHER users and research software engineers wishing to carry out Outreach activities. Over the next few months we will continue to promote this pack and seek feedback. Similarly, the teachers pack is now complete and published online, we will continue to support teachers and seek feedback to enhance these materials.

Impact

The final impact case study has been published, this is on understanding ice formation and nucleation. For all the case studies we have produced postcards, and these have been used at a range of events: SC, ISC, Big Bang Fair, New Scientist Live, etc. These have proved popular and are a valuable resource when explaining the ways ARCHER is used.

9. Embedded CSE (eCSE)

Overview of eCSE Effort



- The eCSE person months awarded up to and including the 12th 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).
- 881 person months have been awarded so far over 90 awarded eCSE projects, meaning an extra 41 person months were awarded at the final call.

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 (see risk list below).
eCSE05	14	8	94	8	8	8	
eCSE06	9	5	47	5	5	4	1 late final report is being pursued.
eCSE07	16	5	49	5	5	4	1 late final report is being pursued.
eCSE08	21	8	88	8	8	3	2 late final reports are being pursued and 3 have recently

							finished with final reports not yet due.
eCSE09	19	5	62	5	3	0	1 project completed early due to a staff member leaving and a final report is being pursued (see risk list below). The other 2 projects completed recently and final reports are not yet due.
eCSE10	13	6	59	6	1	0	1 project has recently finished and the final report is not yet due.
eCSE11	18	6	49	6	1	0	1 project has recently finished and the final report is not yet due.
eCSE12	23	6	41	5	0	0	
Total	201	90	881	89	72	58	

- 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 panel review and this risk will now be removed.
 - 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 project is now complete, and we are awaiting the final report. This final report is now very overdue despite sending several reminders to the PI and there is a risk the PI will not provide this report. We have given the PI a final deadline for sending the report to us and if this is not received we will note the project as having not reported.
 - 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 is now due but has not yet been received.
 - eCSE08-10: there were issues raised by Cambridge University involving the IP and the relationship with the CASTEP group
 - The issues were resolved, and the project went ahead with the final report having now been received. This project is now complete, and the final report has been received. The report has received a favourable panel review and this risk will now be removed.
 - 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 2017
 - The PI confirms 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.

- 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 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 continue to be monitored via regular contact with the PI.
- eCSE10-1: at the time of the previous quarterly report the contract had not yet been agreed
 - This contract has now been agreed and the risk will be removed.
- eCSE10-5: a change of staffing is required
 - We have discussed this with the PI and have agreed the project has been scaled back and re-staffed but will monitor the situation via regular contact with the PI. The unused funds were used to fund eCSE12 projects at the final panel meeting.
- eCSE10-10: at the time of the previous quarterly report the contract had not yet been agreed
 - This contract has now been agreed and the risk will be removed.
- eCSE11-3: at the time of the previous quarterly report the contract had not yet been agreed
 - This contract has now been agreed and the risk will be removed.
- eCSE12-20: the project runs right up until 30/9/2018 – almost the end of the CSE contract
 - The project will be monitored via regular contact with the PI.
- During this Quarter, the following ARCHER webinar was given on a completed eCSE project:
 - 24 January 2018 – eCSE07-16 – “Massively Parallel OpenMP-MPI Implementation of the SPH Code DualSPHysics”, *Athanasios Mokos, University of Manchester*