ARCHER CSE Service Quarterly Report

Quarter 3 2018

1. Document Information and Version History

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

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2018-09-13	Skeleton document	Andy Turner
0.2	2018-09-28	Added Training sections	David Henty
0.3	2018-10-01	Added eCSE sections	Chris Johnson
0.4	2018-10-03	Added Centralised CSE sections	Andy Turner
0.5	2018-10-05	Added Collaborations and Outputs	Andy Turner
0.6	2018-10-09	Reviewed	Alan Simpson
1.0	2018-10-12	Version for EPSRC	Alan Simpson

2. Executive Summary

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

- Centralised CSE Team:
  - In collaboration with GW4 RSE, we organised a Champions and Tier-2 RSE Workshop that ran directly after RSE18 in Birmingham. This event brought together ARCHER Champions, Tier-2 HPC RSEs and other RSEs with an interest in HPC for a series of technical presentations and discussions on how to take the UK HPC RSE community forward. See: https://goo.gl/7u3oA5. The plan is for the ARCHER Champions initiative to be rebranded as UK HPC Champions to provide local points of contact for researchers to help their use of UK HPC resources.
  - We had an initial meeting to exchange experience and coordinate benchmarking effort between the ARCHER CSE team, GW4 RSE and HEC BioSim; this has since been extended to include the Cambridge RSE team. This sharing of expertise complements the wider work to bring the UK HPC RSE community together in a more coordinated way to provide the best support for researchers.
- Training:
  - We delivered 18 days (419 student-days) of face-to-face training in the quarter at 7 different locations, with an average feedback score better than "Very Good".
  - To promote HPC to new users and increase the diversity of science done on the national services, we delivered our first run of the "HPC Carpentry" course. This brand new course takes the interactive, hands-on philosophy of the existing Software Carpentry and Data Carpentry courses and applies it to HPC.
  - In order to enable increased participation, we are running a new online version of our existing OpenMP course starting in October. This will use the same format as the successful online MPI course run in January, delivered using Blackboard Collaborate over four successive Wednesday afternoons.
  - The training plan for 2019 has been drafted and will be circulated to the ARCHER CSE Training Panel in October 2018.
- eCSE:
  - All 90 projects awarded over the first 12 eCSE calls issued have started with 86 of these having now been completed. The remaining 4 projects are all planned to finish by 31 October 2018. Of those completed, 69 final reports have been received.
  - The eCSE13 call was opened on 5 June 2018 and closed on 17 July 2018, receiving 21 proposals. At the panel meeting on 19 September 2018, 10 projects were selected with 96 person months awarded. Projects selected at this call will run between 1 October 2018 and the expected end of the ARCHER service, 18 November 2019.

3. Collaborations and Outputs Summary

- Presentations:
  - *Measuring the Financial Return on Investment of the ARCHER eCSE Programme*, **Lorna Smith**, 4 September 2018, RSE18, University of Birmingham, UK
  - *International Session*, **Andy Turner**, 4 September 2018, RSE18, University of Birmingham, UK
  - *Open Source Benchmarking*, **Andy Turner**, 5 September 2018, Tier-2/Regional HPC RSE Workshop, University of Birmingham, UK
  - *ARCHER Champions*, **Alan Simpson**, 5 September 2018, Tier-2/Champions RSE Workshop, University of Birmingham, UK
  - *Cellular automata beyond 100k cores: MPI vs Fortran coarrays*, **Luis Cebamanos**, 24 September 2018, EuroMPI, Barcelona, Spain
- Meetings:
  - RSE Leaders Meeting, **Andy Turner**, 23 July 2018, University of Birmingham, UK
  - Parallel in time methods discussion and seminar, **Rupert Nash, David Scott**, 9 August 2018, University of Edinburgh, UK
  - Mellanox discussion meeting, **Andy Turner, Alan Simpson**, 22 August 2018, University of Edinburgh, UK
  - Programming Teaching Symposium, **Andy Turner, Mario Antonioletti, Fiona Reid**, 24 August 2018, University of Edinburgh, UK
  - Mellanox discussion meeting, **Andy Turner, Alan Simpson**, 28 September 2018, University of Edinburgh, UK
  - Arm discussion meeting, **Andy Turner, David Henty, Luis Cebemanos, Nick Johnson, Neelofer Banglawala, Arno Proeme,** 28 September 2018, University of Edinburgh, UK

4. Forward Look

- Centralised CSE Team:
  - We plan to incorporate per-job IO statistics provided by the Cray LASSi tool into the SAFE so that users can understand better how their applications use the parallel file system and so improve their efficiency in using ARCHER.
  - We will produce a performance report comparing different HPC compute-node architectures (e.g. Intel Xeon, Nvidia GPU, Cavium ThunderX2 Arm) for the ARCHER benchmarks to help users understand which systems they should choose for their research.
  - We are investigating the use of Linux containers for HPC workflows, which have the potential to make HPC systems more flexible and able to meet the needs of a wider variety of researchers.
- Training:
  - We are continuing our successful collaborations with the Alan Turing Institute by running the HPC Carpentry course at The Turing Institute in London in December.
  - The 2019 training plan will be considered by the ARCHER training panel in October
  - To help prepare the UK community for future technology, we are running two events – a webinar and a face-to-face course – focusing on using Arm-based HPC systems such as the Tier-2 Isambard service and the upcoming Catalyst systems.

5. Contractual Performance Report

This is the contractual performance report for the ARCHER CSE Service for the Reporting Periods: July 2018, August 2018 and September 2018.

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	Jul-18		Aug-18		Sep-18		Q3 2018	
Metric	**Perf.**	**SP**	**Perf.**	**SP**	**Perf.**	**SP**	**Perf.**	**Total**
QE1	100%	-2	100%	-2	100%	-2	100%	-6
QE2	100%	-2	100%	-2	92%	-2	100%	-6
TA1	90%	-1	100%	-1	100%	-1	100%	-3
FB1			100%	-2			100%	-2
Total		-5		-7		-5		-17

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-18		Aug-18		Sep-18		Q3 2018	
Metric	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-18	Aug-18	Sep-18
Total Service Points	-6	-8	-6

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 286 queries were resolved by the CSE service in the reporting period.

Metric	Jul-18	Aug-18	Sep-18	Total
Course Registration	65	81	71	217
eCSE Application	6	6	4	16
In-Depth	7	5	12	24
Course Enquiry	2	0	0	2
Technical Assessment: Grant	6	5	3	14
Technical Assessment: Instant	4	2	1	7
Technical Assessment: RAP	0	1	5	6
Total	**90**	**100**	**96**	**286**

1 query feedback response was received on In-depth queries in the reporting period. This represents a 4% return rate for feedback forms. The 1 response 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	58%
Compilers and system software	4	17%
User programs	3	13%
Disks and resources	1	4%
Performance and scaling	1	4%
Other	1	4%

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.

Q980346: MPI compilation help - NWChem6.8 / ChemShell3.7

An ARCHER user was encountering issues in getting the ChemShell framework to interface properly with the parallel NWChem computational chemistry package. The difficulty arises as both ChemShell and NWChem are complex programs to compile and install, and there is then the additional complexity of ensuring that they are built in such a way that they can be interfaced successfully. This makes installation challenging even for an experienced user. The CSE team was able to work through the complex compile and installation process with the user to allow them to produce a version of ChemShell interfaced with NWChem. Thanks to this help from the ARCHER CSE team, the user was able to use ARCHER successfully for their research. The final, updated

compile process was also passed back to the ChemShell developers for inclusion in their documentation.

Q1051944: CrayPat

The Cray Performance Analysis Tool (CrayPAT) is a useful piece of software installed on ARCHER that allows users to instrument their applications to gain insight into where performance issues may lie to allow them to target code improvements or understand issues with using their software in parallel. An ARCHER user was having issues using CrayPAT on the KNL system with their application as the CrayPAT instrumentation step was resulting in errors without useful error messages. After a detailed investigation, the ARCHER CSE team found that the application was using a number of external libraries that did not have static versions available, there were settings in the user's build process that were causing issues for the CrayPAT instrumentation and the default linker was not working properly with their application. We provided the user with an updated set of instructions and settings to use CrayPAT with their application and this allowed them to proceed with the performance analysis.

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).

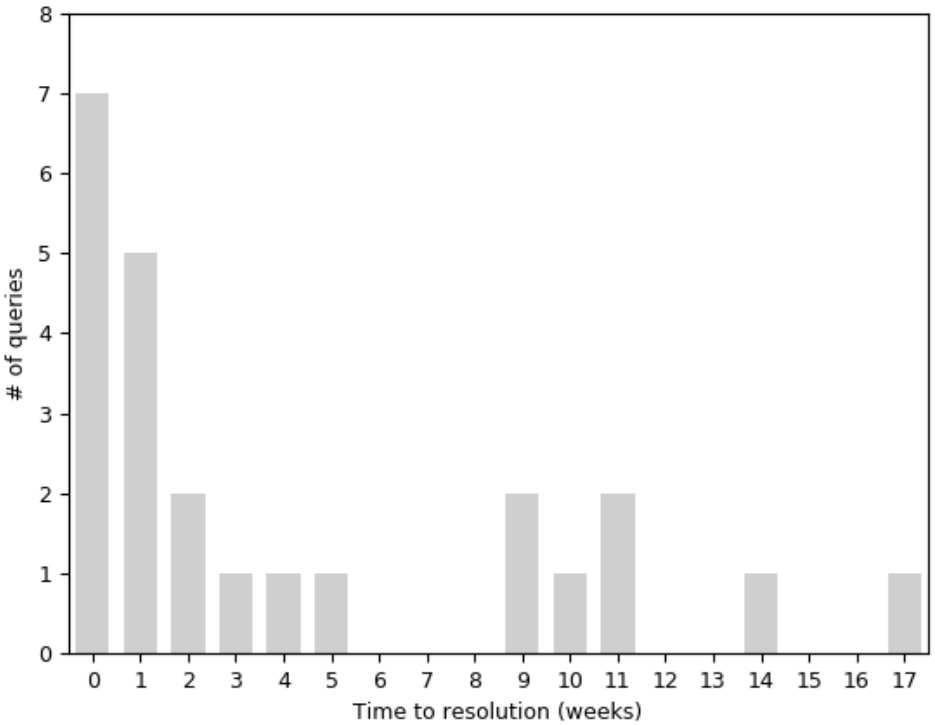

Plot of numbers of In-Depth queries received per quarter:

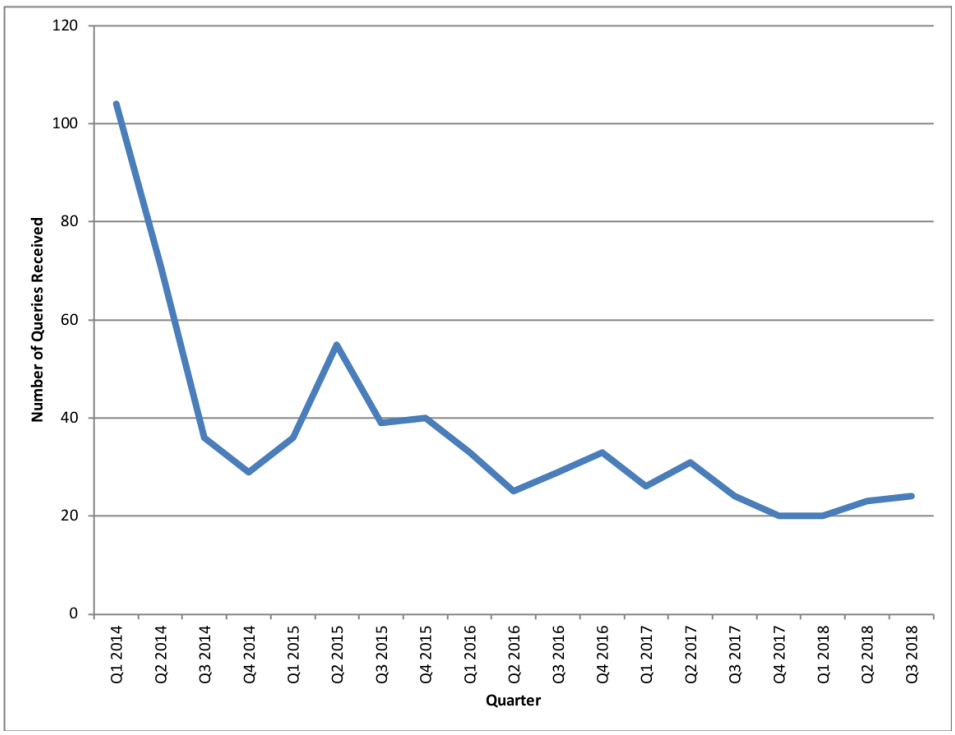

Technical Assessment Analysis

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).

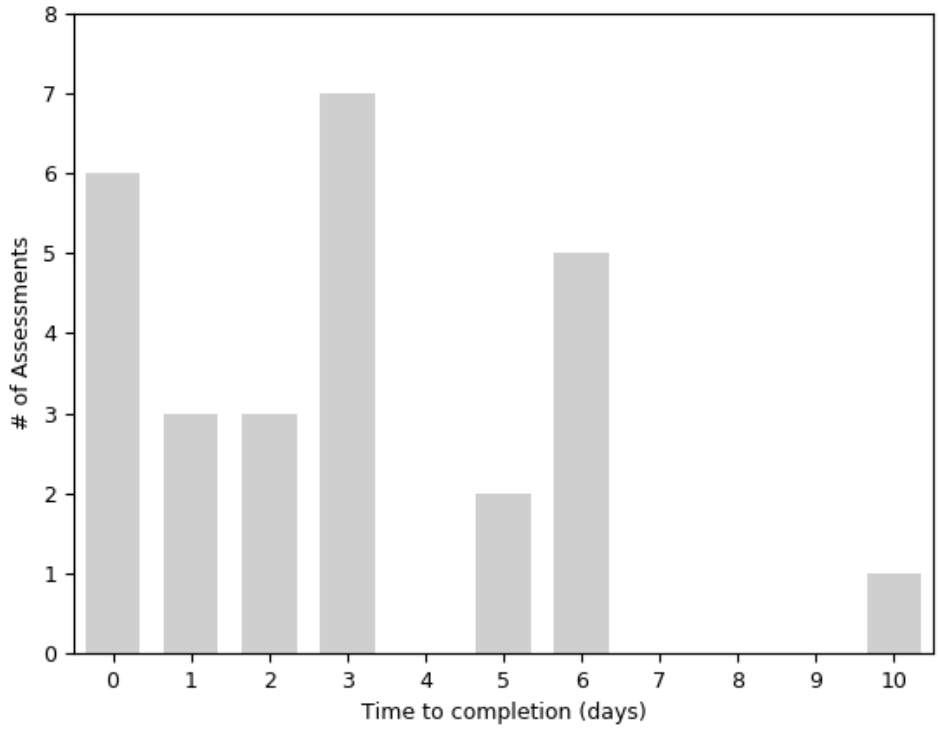

Plot of numbers of Technical Assessments received per quarter:

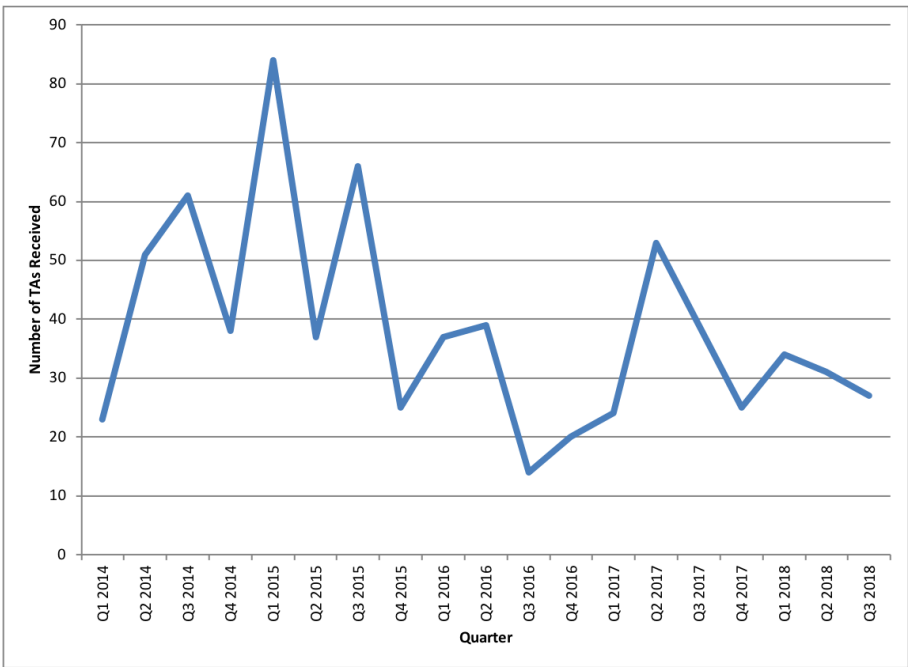

7. Centralised CSE Team: Continual Service Improvement

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

Comparative Benchmarking

Following on from our report comparing the performance of different Intel Xeon HPC systems across the UK we have started the comparisons of different processor and compute node architectures using the same benchmarking set. The architectures investigated so far include:

- Intel Xeon nodes (Ivy Bridge, Broadwell, Skylake Gold, Skylake Silver)
- Nvidia GPU-based nodes (GPUs connected together by PCIe or by NVlink)
- Cavium ThunderX2 Arm64 nodes
- AMD EPYC (Naples) nodes

In the first instance, we will compare the performance of single nodes for these different architectures. This approach is taken for two reasons: firstly, this minimises effects due to different interconnect technologies allowing us to draw clearer conclusions as to the effect of different processor architectures; secondly, the Cavium ThunderX2 Arm64 system is currently restricted to single compute nodes without an HPC interconnect. We plan to publish a report on comparing single node performance between different architectures in Q4 2018.

We also had an initial meeting with the GW4 RSE team and HEC BioSim, who are both also working on comparative benchmarking, to compare approaches and to share experience and results. This meeting has allowed us to coordinate better our work. For example, the ARCHER CSE team have adopted one of the HEC BioSim GROMACS benchmarks to allow for a wider comparison of GROMACS performance across systems than would otherwise have been possible. We plan to put in place a regular schedule of meetings for RSEs across the UK interested in comparative benchmarking to share experience going forward.

Containers for HPC

There is a huge amount of interest in the worldwide HPC community in the use of Linux containers to allow for more flexible use of HPC facilities and to potentially allow compute environments to be more reproducible across different technologies. Containers have been widely adopted in High Throughput Computing (HTC) use through technologies such as Docker (https://www.docker.com/), but many of these approaches are not appropriate for use in HPC environments. One promising technology for using containers on multi-user HPC systems is Singularity (https://www.sylabs.io/). We have been investigating the use of Singularity over the past year and have found it useful for running serial workflows on HPC systems. We have recently begun more detailed investigations into using Singularity in parallel across multiple nodes with applications using MPI. We plan to use the ARCHER benchmarks to test the functionality and performance of Singularity containers for different HPC use cases. This work will form the basis of a number of dissemination activities on using containers in HPC, including: an ARCHER white paper and an ARCHER webinar. We also plan to propose a tutorial on containers in HPC to ISC19 in collaboration with the Pawsey Supercomputing Center (Perth, Australia).

8. Training

This quarter, the CSE Service has provided a total of 18 days (419 student-days) of face-to-face training across 7 different locations and 1.0 days of interactive online tutorials (average attendance 15 per tutorial).

Month	Dates	Course	Location	Days	Attendee
Jul 2018	2-3	Hands-on Intro to HPC	Edinburgh	2	33
	4-6	Message-Passing Programming	Edinburgh	3	30
	25	Advanced OpenMP	Cambridge	3	17
	17-19	Version Control	Online	0.5	
	25-26	HPC Carpentry	Leeds	2	32
	30-31	Advanced MPI	Milton Keynes	2	10
Sep 2018	12	eCSE Benefits Realisation	Online	0.5	
	13-14	Shared-Memory Programming with OpenMP	London	2	13
	17-18	Scientific Python	Durham	2	30
	20-21	Data Carpentry	Swansea	2	21

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 88 feedback forms, a response rate of 47%.

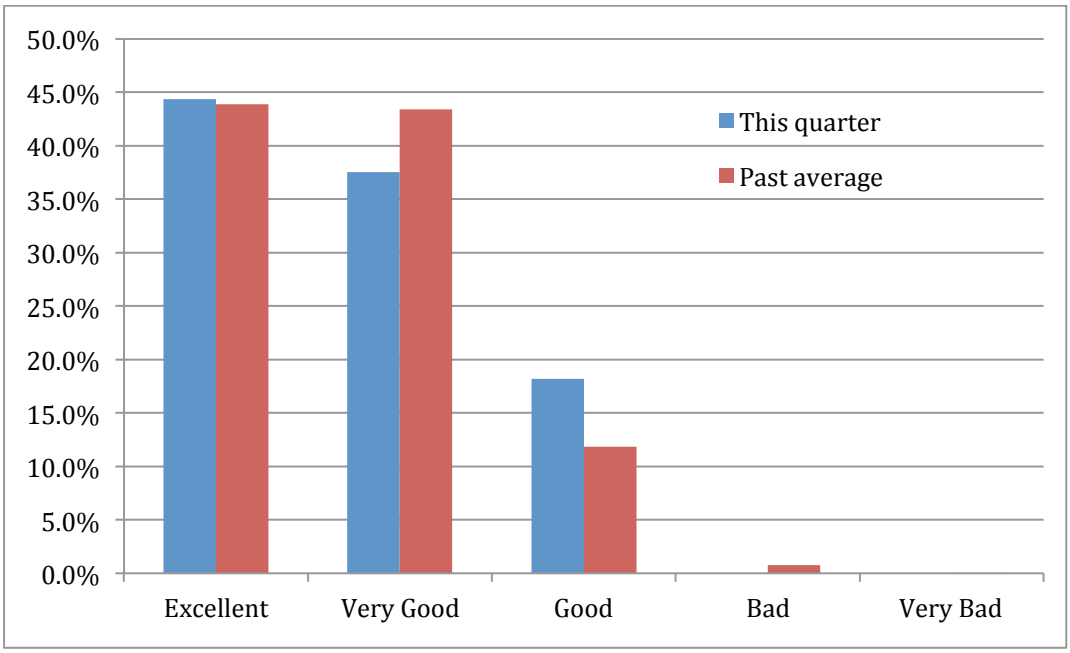

We reported in the previous report that we had received one overall score of "Bad" in Q2. The attendee entered their personal information and stated that they were happy to be contacted. We contacted them to ask for more details but did not receive any response.

15 days of face-to-face training are planned for the third quarter of 2018, plus 3.0 days online.

Month	Dates	Course	Location	Days	Attend
Oct 2018	9	Hands-on Intro to HPC	Edinburgh	1*	
	17	The TPLS Solver and PETsC	Online	0.5	
	24	OpenMP online (four consecutive Wednesday afternoons)	Online	2	
Nov 2018	6-8	Parallel Design Patterns	Oxford	3	
	26-27	GPU Programming with CUDA	London	2	
	28-30	Message-Passing Programming with MPI	Newcastle	3	
Dec 2018	3-4	Programming the ARM CPU	Edinburgh	2	
	6-7	HPC Carpentry	London	2	
	6-7	Scientific Python	Newcastle	2	
	12	Isambard: The UK's Tier-2 ARM-based Supercomputer	Online	0.5	

*This course also forms part of The University of Edinburgh's celebration of Ada Lovelace Day 2018

9. Embedded CSE (eCSE)

Overview of eCSE effort

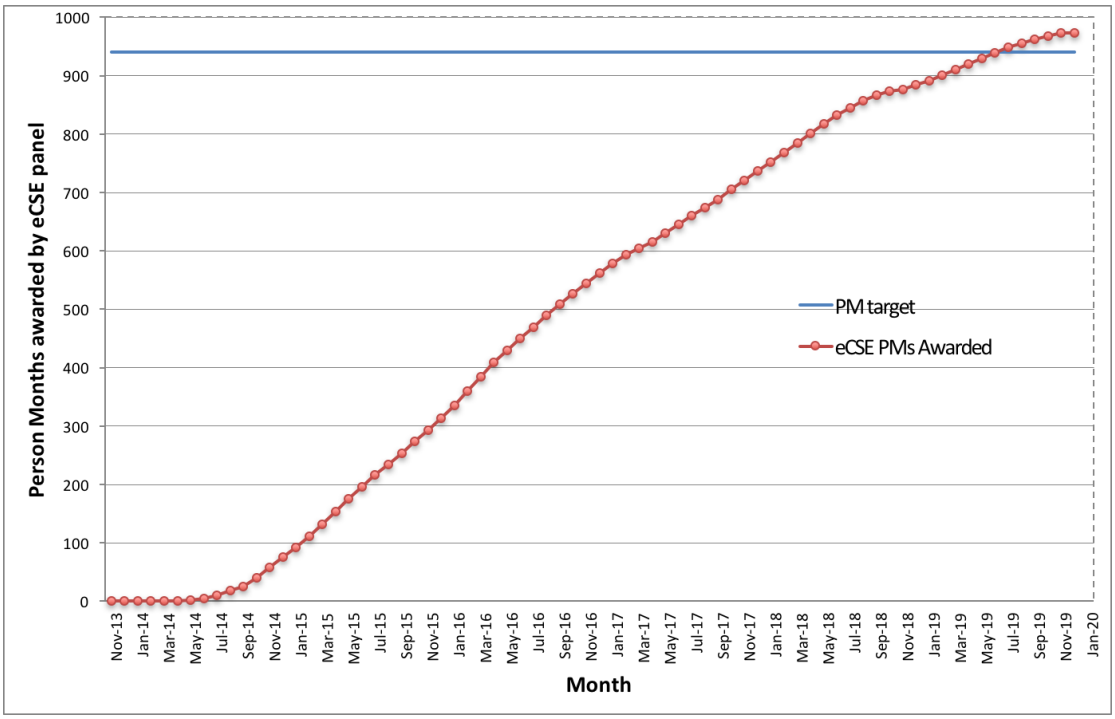

- The eCSE person months awarded up to and including the 13th eCSE call are shown in red.
- We committed to awarding at least 941 person months by the end of the project (14 FTEs for 5 years, and 8.4 FTE for year 6).
- 973 person months have been awarded across 100 eCSE projects meaning an extra 32 person 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	9	1 late final report is being pursued.
eCSE04	16	8	82	8	8	8	
eCSE05	14	8	94	8	8	8	
eCSE06	9	5	47	5	5	5	
eCSE07	16	5	49	5	5	5	
eCSE08	21	8	88	8	8	6	2 late final reports are being pursued.
eCSE09	19	5	58	5	5	3	1 project completed early due to a staff member leaving and a final report is being pursued (see risk list below). The final report for the other completed project is due shortly.
eCSE10	13	6	59	6	5	1	4 final reports are

							due during the next quarter.
eCSE11	18	6	49	6	5	1	2 late final reports are being pursued. The other 2 are due during the next quarter.
eCSE12	23	6	41	6	4	0	
eCSE13	21	10	96	0	0	0	
Total	**222**	**100**	**973**	**90**	**86**	**69**	

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