ARCHER2 Quarterly Report

April – June 2022

EPCC

The University of Edinburgh

1. Document Information and Version History

Version:	1.0
Status	Release
Author(s):	Lorna Smith, Juan Rodriguez Herrera, Chris Johnson, Xu Guo, Anne Whiting, George Beckett
Reviewer(s)	Alan Simpson

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2022-06-08	Initial draft	Lorna Smith
0.2	2022-06-28	Adding ISO, feedback and impact data	Anne Whiting
0.3	2022-07-04	Adding ARCHER2 CSE queries performance report, statistics and analysis	Xu Guo
0.4	2022-07-13	Complete draft	Lorna Smith
0.5	2022-07-14	Reviewed	Alan Simpson
1.0	2022-07-15	Version for UKRI	Lorna Smith, Alan Simpson

ARCHER 2 Quarterly Report

This section of the report covers the period April 2022 – June 2022 for the ARCHER2 CSE service.

2. ARCHER2 Executive Summary

- A series of short projects have been started, looking to enhance the overall environment for users. This includes a joint activity of CSE and the HPE Centre of Excellence team to develop a lightweight API for gathering power performance statistics generated by parallel codes.
- In addition, work to develop a package-management tool, based on Spack, has resumed following improvements to Cray Programming Environment-Spack compatibility. This will help streamline the process of package management.
- CSE has access to an HPE-internal GPU system and is gaining knowledge of the platform in preparation for the introduction of a small GPU resource on ARCHER2.
- CSE presented three papers based on ARCHER2 work to Cray User Group meeting, in May. This allowed the exchange of best practice with sites running similar systems.
- The team had a booth at the Big Bang Fair in Birmingham in June, which was very well attended by school children from a wide geographical area. This allowed the team to explain the importance of supercomputing to society and raise awareness around career opportunities in computational science.
- A total of 21 days of training, including courses and virtual tutorials, have been delivered using the full ARCHER2 system. These courses covered a wide spectrum of user interests, including courses on key packages such as GROMACS and Code_Saturne as well as an Advanced MPI course and a course looking at making efficient use of ARCHER2.
- The 7th eCSE call is complete, with a range of projects funded. The panel meeting involved a group of early career researchers present as observers to gain experience of the panel and review process.
- With the aim of reducing barriers to applicants from less traditional areas, a blog article has been written to describe the opportunities presented by the eCSE programme to PDRAs and RSEs. An eCSE webinar followed by an extended tutorial session were presented while the eCSE07 call was open.
- An audience-led panel has been accepted to RSECon22, taking place in Newcastle on 5-8 September 2022. This panel has a focus on highlighting opportunities for RSEs and barriers surrounding accessing software engineering funding.
- EPCC passed the stage 1 external audit of ISO22301, the business continuity standard. Our commitment here recognises the importance of making sure that user access to services and to their data is maintained.

3. ARCHER2 Forward Look

- The CSE team will be well represented at RSECon in September, with an aim of highlighting the value and opportunities available to the RSE community from ARCHER2. An audience-led panel has been accepted. Led by the eCSE team, this has a focus on highlighting opportunities for RSEs and barriers surrounding accessing software engineering funding. The CSE team will also contribute to our booth, aiming to highlight support options for current and potential users.
- The CSE team will contribute to the roll-out of the ARCHER2 Gazelle (solid-state) storage in the coming period, subject to access to the ARCHER2 Test and Development system. Activities will include benchmarking and developing guidance and documentation on how best to utilise this storage.
- The CSI projects will continue, with the team looking at performance monitoring tools, preparation for the arrival of the small GPU component, and investigation of package management tools.
- A set of upcoming virtual webinars will introduce several topics that were submitted and presented at CUG22. This will allow this information to be disseminated to the wider ARCHER2 community.
- We continue to test and improve our service delivery, information security and business continuity readiness. We have our ISO 9001 Quality and ISO 27001 Information Security external audits in September and also the renewal of our Cyber Essentials certification.
- Planning is underway for New Scientist Live, a large outreach event in London during October. This provides an opportunity to showcase supercomputing, both to school children and to the wider public.
- Having recently completed the seventh call for the eCSE programme, a number of the early projects are now complete and we have begun to receive final reports. During this period, we will be looking to publicise the outcomes and achievements of these projects.
- Having completed 6 months of the full service, an initial set of benefits realisation data has been produced and will be refined in conjunction with UKRI to meet the needs of the funding councils.

4. ARCHER2 Centralised CSE Team

The CSE team has continued to support users to undertake their computational research during the period without significant exceptions.

A particular highlight of the period was a strong representation by the CSE (and other ARCHER2 teams) at the Cray User Group meeting, held in Monterey Bay, California. Two papers were submitted and accepted from the core team (David Henty and Holly Judge), along with a third from an EPCC-based ARCHER2 user (Nick Brown):

- Performance of Parallel IO on the 5860-node HPE Cray EX System ARCHER2, David Henty
- Efficient use of MPI+OpenMP on a Cray EX supercomputer, Holly Judge
- Predicting batch queue job wait times for informed scheduling of urgent HPC workloads, Nick Brown

Juan Rodriguez Herrera also participated as an invited member of the programme committee for the meeting. This meeting was particularly useful to allow sharing of best practice and experience between sites running similar systems.

The work presented at CUG will be re-run for the UK HPC community, via the ARCHER2 webinar series, over the coming weeks.

Andy Turner represented the CSE team at the International Conference on Computational Science, 23 June 2022 (online/Brunel University), presenting a talk titled “Exascale Computing and UQ, SEAVEA VVUQ on the Exascale”.

Continual Service Improvement (CSI) Projects

The full service has been operational for just over 6 months. The CSE team continue to support users through in-depth queries but, as the system has settled into normal production, the team has been able to look to carry out a series of small improvement projects. These projects look to enhance the ARCHER2 environment for users, through best practice guides, enhanced tools and improved documentation.

The following CSI activities have been carried out over the last quarter. In addition to these activities, the next quarter will also see a focus on the use of the ARCHER2 Gazelle (solid-state) storage as it is rolled out in the coming period, subject to access to the ARCHER2 Test and Development system.

Porting and Optimising CP2K for the ARCHER2 Main System

Holly Judge’s has completed a best practice guide on building and using CP2K on ARCHER2, in collaboration with the BioExcel project. The best-practice guide has been published on the web (https://docs.bioexcel.eu/qmmm_bpg/en/main/running_cp2k/running_cp2k.html) and also published as a paper (https://doi.org/10.5281/zenodo.6591574). An ARCHER2 blog article will be produced highlighting the value of the work to the ARCHER2 community.

Monitoring the ARCHER2 Power Management Hardware Counters

This project is a joint activity of CSE and the HPE Centre of Excellence team to develop a lightweight API for gathering power performance statistics generated by parallel codes. Having access to energy usage data for ARCHER2 jobs is expected to be extremely useful in understanding and improving power efficiency, plus be a useful input to – for example – net-zero targets.

This tool was previously deployed on ARCHER and HPE have developed an ARCHER2 port of the baseline toolset. This will be extended by CSE to capture and report energy usage within MPI applications.

ARCHER2 GPU Platform Evaluation

Before the ARCHER2 GPU nodes are installed, HPE has provided the CSE team with access to a comparable internal system, to allow the CSE team to do an initial evaluation of the programming environment, pursue code-porting strategies, and start to define the service offering for ARCHER2.

The team (led by Michael Bareford) has made good progress during May and June, successfully building small-scale HIP-, OpenMP-, and Kokkhos-based applications with the Cray Programming Environment. They have also tested the built-in debugging and profiling tools.

Confident performance testing requires access to a GPU-aware version of the MPI library. This is not currently available on the evaluation system, though HPE staff have agreed to install a suitable library (specifically ROCm, Version 5) soon.

The team is documenting their findings in a metrics paper, which follows a similar structure to that used for CSE testing of ARCHER2 during the phased installation.

Package Management with Spack

In 2020, the CSE team evaluated several package-management systems (namely, EasyBuild and Spack) to streamline and improve package management on ARCHER2. At the time, incompatibilities with the HPE EX platform meant neither option was suitable, though the CSE team began monitoring developments in the HPC community, with an interest in revisiting the situation once the HPE EX platform was better established.

Two years on, both HPE and the Spack developers have made changes to better support the integration of Spack with the ARCHER2 module system. Prompted by this, the CSE team has resumed the evaluation.

Initial work has begun to define Spack as a service on ARCHER2, with a two-tier usage model: one tier would be for CSE, to install software centrally, with the software and associated modules generated by Spack available to all users. The second tier of usage would allow users themselves to install software in their own work directories. This might constitute software packages too obscure for central installation and support, using different versions, compilers, or libraries, or in unusual configurations.

Use of Spack would, if successful, allow users this extra flexibility. The CSE service in turn could rely on the built-in packages or, if necessary, write their own. Once a package's given build specification is known to complete successfully, this should be easily repeatable in the future; for example, when upgrading to new CPE versions, for regression testing, or in the case of disaster recovery.

Initial testing shows that updates to Spack allow it to work well with the new CPE versions for HPE EX systems. The very latest CPE versions also provide software manifests to allow overlaying of Spack onto pre-existing CPE software.

Moving forward, the CSE team will determine how to set up the two-tier use model, noted above, and establish a process for the management (via Spack) of licensed software, such as CASTEP and VASP.

5. ARCHER2 Performance Report

This is the performance report for the ARCHER2 CSE Service for the Reporting Periods from April 2022 – June 2022.

The metrics were specified by EPSRC in Schedule 2.2 of ARCHER2 CSE Service Contract.

CSE Query Metrics

- **ARCHER2_CSE_Level1 (MTR)**: The Median Time to Resolution, as measured by Working Days (WDs), of all CSE queries falling within Level 1 resolved by the Contractor in the Reporting Period. *MTR applicable to OY3: Service Threshold: >4.2 WD; Operating Service Level: >1.2 WD, ≤2.2 WD.*
- **ARCHER2_CSE_Level2 (MTR)**: The Median Time to Resolution, as measured by Working Days (WD), of all CSE queries falling within Level 2 resolved by the Contractor in the Reporting Period. *MTR applicable to OY3: Service Threshold: >26 Working Days (WD); Operating Service Level: >11 WD, ≤16 WD.*
- **ARCHER2_CSE_Level3 (MTR)**: The Median Time to Resolution, as measured by Working Days (WD), of all CSE queries falling within Level 3 resolved by the Contractor in the Reporting Period. *MTR applicable to OY3: Service Threshold: >57 Working Days (WD); Operating Service Level: >27 WD, ≤37 WD.*
- **ARCHER2_CSE_TA (%):** The percentage of the total number of Technical Assessments (TAs) assigned to the Contractor in the Reporting Period completed prior to the commencement of the applicable TA Target Completion Date after the assignment of such Technical Assessment to the Contractor*. TA Target Completion Date in OY3: 7 WD; Service Threshold: <90.00%; Operating Service Level: 95.00-97.49%.*
- **Initial Response to Queries (%):** The percentage of the total number of CSE queries assigned to the Contractor in the Reporting Period responded to within 3 Working Hours. *Service Threshold: <96.00%; Operating Service Level: 98.00 – 98.99%.*
- **Query User Satisfaction (%):** The percentage of the total number of query satisfaction surveys completed in each Reporting Period, rating the quality of the resolution of Queries by the Contractor as “Good”, “Very Good” or “Excellent”. *Operating Service Level: 82.00 – 87.99%.*
- **Training User Satisfaction (%):** The percentage of all training satisfaction surveys completed in each Service Period, rating the Contractor as “Good”, “Very Good” or “Excellent”. *Operating Service Level: 88.00%-92.99%.*

Metric	April 2022		May 2022		June 2022		Q2 2022	
	Perf	Points	Perf	Points	Perf	Points	Perf	Points
ARCHER2_CSE_Level1 (MTR)	0.1WD	-2	0.1WD	-2	0.1WD	-2	0.1WD	-6
ARCHER2_CSE_Level2 (MTR)	0.2WD	-2	0.5WD	-2	0.4WD	-2	0.4WD	-6
ARCHER2_CSE_Level3 (MTR)	21WD	-0.5	-		-		21WD	-0.5
ARCHER2_CSE_TA (%)	100%	-1	100%	-1	100%	-1	100%	-3
Initial Response to Queries (%)	100%	-1	100%	-1	100%	-1	100%	-3
Query User Satisfaction (%)	97.3%	-2	98.2%	-2	100%	-2	98.4%	-6
Training Satisfaction (%)	100%	-1	100%	-1	100%	-1	100%	-3
Total		-9.5		-9		-9		-27.5

125 query feedback responses were received on query resolution in the Reporting Period. 98.4% of responses had a score of “Good”, “Very Good” or “Excellent”.

6. ARCHER2 CSE Queries

This section provides details on ARCHER2 CSE queries during the Reporting Periods from April 2022 – June 2022.

CSE Query Statistics

The metrics were specified by EPSRC in Schedule 2.2 of ARCHER2 CSE Service Contract.

- **Assigned:** The number of CSE queries assigned to the Contractor within each query resolution category in the Reporting Period.
- **Resolved:** The number of CSE queries resolved by the Contractor within each query resolution category in the Reporting Period.
- **Backlog:** The number of CSE queries assigned to the Contractor that remained unsolved within each query resolution category in the Reporting Period
- **Correspondence:** The average number of pieces of correspondence generated for CSE queries in each query resolution category.
- **First Response:** The average time taken for the Contractor to first respond to the Originator of the CSE query.

April 2022					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
ARCHER2_CSE_Level1	114	114	1	3	0.2hr
ARCHER2_CSE_Level2	61	57	32	12	0.3hr
ARCHER2_CSE_Level3	0	2	3	32	0.2hr
ARCHER2_CSE_TA	3	1	2	11	0.2hr
May 2022					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
ARCHER2_CSE_Level1	61	62	0	3	0.3hr
ARCHER2_CSE_Level2	85	91	26	12	0.3hr
ARCHER2_CSE_Level3	0	0	3	0	-
ARCHER2_CSE_TA	8	4	6	11	0.2hr
June 2022					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
ARCHER2_CSE_Level1	65	65	0	3	0.3hr
ARCHER2_CSE_Level2	57	52	31	11	0.3hr
ARCHER2_CSE_Level3	1	0	4	0	-
ARCHER2_CSE_TA	18	20	4	13	0.2hr
Q2 2022					
Service level	**Assigned**	**Resolved**	**Backlog**	**Correspondence**	**First Response**
ARCHER2_CSE_Level1	240	241	0	3	0.2hr
ARCHER2_CSE_Level2	203	200	31	12	0.3hr
ARCHER2_CSE_Level3	1	2	4	32	0.2hr
ARCHER2_CSE_TA	29	25	4	12	0.2hr

CSE Query Categories

A total of 468 queries were resolved by the ARCHER2 CSE service in the Reporting Period. Resolved CSE queries in the Reporting Period fell into the following categories:

Service level	Category	Number resolved	% Queries
ARCHER2_CSE_Level1	Courses	241	51.5%
ARCHER2_CSE_Level2	3rd Party Software	88	18.8%
	eCSE Applications/Calls	25	5.3%
	Batch system and queues	24	5.1%
	Compilers and system software	13	2.8%
	Login, passwords and ssh	11	2.4%
	Courses	7	1.5%
	Porting, performance and scaling	7	1.5%
	Software errors	7	1.5%
	Software installation	7	1.5%
	Hardware issue	4	0.9%
	Access to services	2	0.4%
	Storage and compute resources	2	0.4%
	Other	1	0.2%
	User behaviour	1	0.2%
	Website and documentation	1	0.2%
ARCHER2_CSE_Level3	3rd Party Software	1	0.2%
	Porting, performance and scaling	1	0.2%
ARCHER2_CSE_TA	UKRI HEC Consortia Call	10	2.1%
	UKRI Grant	7	1.5%
	Pump Priming	4	0.9%
	EPSRC Responsive Mode	2	0.4%
	EPSRC Fellowship	1	0.2%
	Director Time	1	0.2%
Total		468	100.0%

7. ARCHER2 Training

As part of ARCHER2, the service has been developing and delivering online a training programme for the ARCHER2 community. During the second quarter of 2022, the CSE service has provided a total of 21 days of online training, scheduled as follows:

Dates	Course	Days	Attend
5-7 Apr 2022	Efficient use of the HPE Cray EX Supercomputer ARCHER2	3	24
19-20 Apr 2022	Introduction to ARCHER2 for Software Developers	2	7
20 Apr 2022	High-resolution prediction of wind profile and wind turbine aerodynamics using CFD and HPC	0.5	42
21 Apr 2022	Introduction to GROMACS	1	11
22 Apr 2022	ARCHER2 for Data Scientists	1	11
26-27 Apr 2022	Introduction to Modern Fortran	2	8
26 Apr 2022	Top 10* tips for using HPC	0.5	18
27 Apr 2022	HPC-Europa3 project – assessing the nature and role of band alignment of copper oxides heterostructures	0.5	19
27-28 Apr 2022	HPC for CFD using Code_Saturne	1	13
28 Apr 2022	eCSE Webinar	0.5	4
28 Apr 2022	Plotting and Programming with Python	1	24
9-10 May 2022	Performance Analysis Workshop (Scalasca)	2	22
10 May 2022	eCSE Tutorial - Preparing an eCSE Proposal	0.5	5
18 May 2022	Dynamic hydrogen-bond networks for pH sensing at the membrane interface	0.5	10
25 May 2022	A CFD tool for the simulation of renewable energy devices using GPU accelerated hardware	0.5	21
15 Jun 2022	Modelling Cake Formation in a Centrifugal Filter using the Discrete Element Method on ARCHER2	0.5	20
20-21 Jun 2022	Reproducible computational environments using containers	2	13
29-30 Jun 2022	Advanced MPI	2	15

The number of courses and virtual tutorials remains high this quarter, following the momentum of the previous quarter. Several virtual tutorials have been delivered by former HPC-Europa3 visitors, who talked about their experience and progress using an HPC machine such as ARCHER2 to carry out their research.

On the feedback for online 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 4.63, i.e., better than "Very Good". Users provided 49 responses, a response rate of 33%.

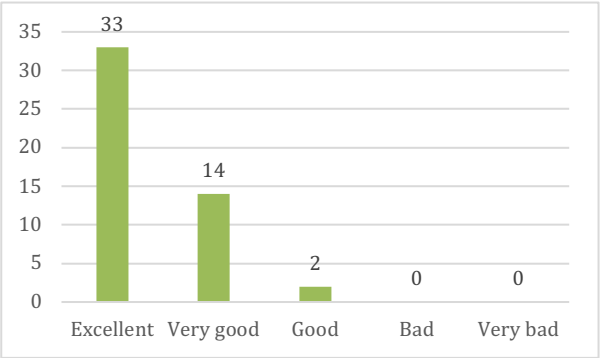

8. ARCHER2 Embedded CSE Programme (eCSE)

eCSE calls 1-7

- eCSE07 has recently completed and award letters are currently being sent out. From the first 6 calls, 414 PMs have been awarded across 45 projects. These are detailed in the table below.
- From call 5 onwards, only proposals with software within the EPSRC remit have been eligible.
- The seventh ARCHER2 eCSE call opened on 19 April 2022 and closed for technical evaluations on 24 May 2022 receiving 13 technical evaluation documents. The call then closed 14 June 2022 receiving 10 proposals. The panel meeting took place on the 12 June 2022 and 4 early career observers attended. Any PIs who submitted a document for technical evaluation but did not go on to submit a full eCSE proposal have been contacted to discuss any possible future applications.

eCSE call	Call Dates	# Technical Evaluations Received	# Proposals Received (EPSRC,NERC)	# PM required (EPSRC, NERC)	# Proposals accepted (EPSRC, NERC)	# PMs Awarded (EPSRC, NERC)
eCSE01	19/05/20 - 07/07/20	25	25 (25,0)	235 (235,0)	13 (13,0)	132 (132,0)
eCSE02	08/09/20 - 27/10/20	13	12 (9,3)	107 (87,20)	7 (4,3)	53 (33,20)
eCSE03	08/12/20 - 16/03/21	15	14 (10,4)	136 (105,31)	8 (6,2)	73 (54,19)
eCSE04	20/04/21 - 08/06/21	13	12 (9,3)	109 (83,26)	7 (4,3)	60 (37,23)
eCSE05	07/09/21 - 26/10/21	10	9 (9,0)	85 (85,0)	5 (5,0)	47 (47,0)
eCSE06	09/12/21 - 15/03/22	7	6 (6,0)	61 (61,0)	5 (5,0)	49 (49,0)
eCSE07	19/04/21 - 14/05/22	13	**10 (10,0)**	**77 (77,0)**	**TBD**	**TBD**
Total		**96**	**88 (78,10)**	**850 (733,77)**	**45 (37,8)**	**414 (352,62)**

The graph below shows the current person months awarded to eCSE projects to date (blue line) along with the number to be awarded for the first 4 years of ARCHER2 (orange line).

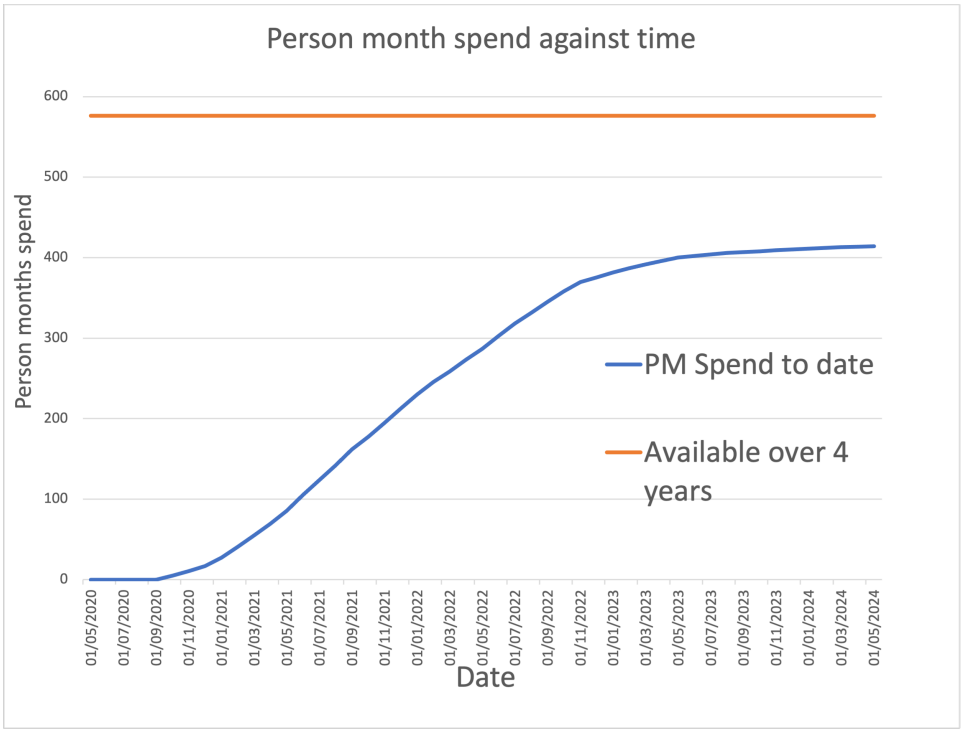
Person month spend against time
Person months spend
600
500
400
300
200
100
0
PM Spend to date
Available over 4 years
01/05/2020
01/07/2020
01/09/2020
01/11/2020
01/01/2021
01/03/2021
01/05/2021
01/07/2021
01/09/2021
01/11/2021
01/01/2022
01/03/2022
01/05/2022
01/07/2022
01/09/2022
01/11/2022
01/01/2023
01/03/2023
01/05/2023
01/07/2023
01/09/2023
01/11/2023
01/01/2024
01/03/2024
01/05/2024
Date

9. ARCHER2 Community Engagement, Outreach, Collaboration and Impact

Benefits Realisation

We have been working with UKRI to develop reports which will help them demonstrate the business case for future investment in the HPC services which delivers so much exciting and innovate science. When users create accounts in the SAFE and on ARCHER2, they add data about themselves, and we collect data on the work run on ARCHER2 such as jobs, disk space, codes used. This data is primarily used to run the service and to ensure each project has access to the resources awarded to it. Summarised statistical data is also used to contribute to what is referred to as benefits data for UKRI for their business cases. The privacy statements on the ARCHER2 website explain how such data is used and why https://www.archer2.ac.uk/about/policies/.

Blogs

The team releases a regular series of blog, providing users and the wider community with information on the activities of the service and on areas of potential interest. Eight blogs have been published this quarter on a range of topics:

- Disaster recovery tests
- ARCHER2 driving test launch
- Software usage data
- Computing Insite UK
- Commercial access to ARCHER2 and Cirrus
- Business continuity management
- Getting started using ARCHER2 and HPC

We would be really keen to hear from any of our users who would like to publicise their projects and the science delivered either in the form of a blog or a case study.

Community and Outreach Activities

The team have been actively engaging in education and outreach activities over the last quarter.

A large team of staff presented a booth at the Big Bang Fair in Birmingham on the X-Y June 2022. The event had over 30,000+ attendees and our booth was very busy across the three days. The team presented activities aimed at highlighting the importance of supercomputing to society and raising awareness of careers in computational science.

Holly Judge attended the BioExcel Summer School. which ran in Sardinia in week beginning 14th June, with school attendees using ARCHER2 for the tutorial sessions. See:
https://bioexcel.eu/events/bioexcel-summer-school-on-biomolecular-simulations-2022

CSE staff contributed to a successful run of the International HPC Summer School (IHPCSS) in Athens during 19th—24th June. David Henty taught the classical MPI and OpenMP track, Weronika Filinger ran the mentoring programme, and Ludovic Capelli (EPCC) ran the programming challenge. During the event, the participants also had an opportunity to participate in focus groups for women and underrepresented groups.

Diversity and Inclusivity

The CSE team supported the Women in HPC workshop and co-located events at ISC'22 in Hamburg. The workshop focused on the topic of creating a diverse and inclusive community (summarised in this blog article - https://womeninhpc.org/women-in-hpc-events/whpc-heads-to-isc22). Preliminary results from survey data show that the event was very well received. The WHPC posters and poster reception organised on the ISC diversity day garnered a lot of interest as well.

The SIGHPC Education chapter, with the support from the ACM, launched a new award for Outstanding Contributions to Computational Science Education (https://sighpceducation.acm.org/edaward.html). One of the important criteria is the ability to engage, serve and retain a diverse community of participants. Weronika Filinger is part of the award committee, and the deadline for applications is 1st July 2022.

Quality Management, Information Security and Business Continuity

With three ISO standards in place or preparing to be in place, we have an annual cycle of improvement work and internal audits to ensure we are maximising the benefits of this work and the focus it brings to the services we deliver for our users and for their data that we store for them. The external audits in the autumn give a focus to these activities, but we are keen to continue to actively use the tools they provide to ensure we use best practice. We hope that all our users are happy with the services they receive and are always pleased to receive feedback which helps us identify areas where more work is required.

We recently passed the stage 1 external audit of ISO22301, the business continuity standard. EPCC recognises the importance of making sure that user access to services and to their data is maintained. We have therefore been planning and testing to ensure that our processes and plans are in place and working to help prevent issues and to deal with any that do arise in a timely manner. The stage 1 audit is an initial inspection of our written plans and processes by an accredited external auditor. The auditor then returns in October to carry out a full audit to ensure that we are carrying out the tests and improvements promised.