



ARCHER2 Quarterly Report

July– September 2023

EPCC

The University of Edinburgh



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	2023-09-26	Initial draft	Lorna Smith
0.2	2023-09-29	Added ISO, benefits, blogs data	Anne Whiting
0.3	2023-10-02	ARCHER2 CSE queries performance report, statistics and analysis added	Xu Guo
0.4	2023-10-03	Added eCSE	Chris Johnson
0.5	2023-10-08	First full version	Lorna Smith
0.6	2023-10-10	Reviewed	Alan Simpson
1.0	2023-10-10	Version for UKRI	Alan Simpson, Lorna Smith

ARCHER 2 Quarterly Report

This section of the report covers the period July 2023 – September 2023 for the ARCHER2 service.

ARCHER2 Executive Summary

- The CSE team have been working closely with HPE this quarter, providing support towards diagnosing and resolving a memory leak that caused degraded job reliability and node availability issues during July. Further, the team has supported users and contributed to the diagnosis and testing of a bug in the Lustre software which has been impacting different user groups to varying degrees.
- The CSE has added IO500 measures to the Reframe suite, complementing the bandwidth data gathered from BenchIO. This suite allows CSE to measure the performance of the work file-system.
- The CSE team was involved in ensuring that EPCC passed our ISO 9001 (quality service delivery), ISO 27001 (information security) and ISO 22301 (business continuity and disaster recovery) annual external audits, with all aspects of the CSE service audited.
- The eleventh eCSE call (eCSE11) opened on 12 September 2023 with the scope of the call widened to allow proposals for projects that will prepare software for a future Exascale successor. Projects involving GPU development can get access to the new GPU development cluster.
- A total of 8.5 days of online and face-to-face training have been delivered this third quarter of the year.
- During the summer, the team hosted three research placement students for 3 weeks as part of the Nuffield research placement programme. This gives the students the opportunity to experience working in a software development environment as they look to develop their careers post school.
- The CSE team has begun work to revise the application benchmark suite, previously developed during the ARCHER service to reflect topical requirements of the UK HPC community and likely future ambitions over the next few years.

ARCHER2 Forward Look

- During this quarter, the CSE team will prepare for the ARCHER2 Science Days event that will take place in March 2024 and will showcase and celebrate the science being achieved on ARCHER2.
- The team will continue to develop the application benchmark suite with the aim for an outline of the required suite to be ready for internal review within CSE by the end of October, and that a first public version should be available by the end of the year.
- The team are working on a plan to deliver a series of Capability Days on the service, coupled with Capability Use workshops. If approved, the aim is to provide users with the opportunity to test and explore capability jobs on a regular basis.
- The small GPU component of ARCHER2 is due to arrive during this quarter or early 2024. The CSE team will therefore look to prepare for the arrival of this component, developing user documentation and training for this system.
- The team will continue to develop collaboration opportunities with the expanded CoE team around areas such as consortium engagement, monitoring power and energy, network and IO, capability usage, and new technologies.
- EPCC will attend New Scientist Live at the ExCeL in London during October, with a booth on the main exhibition floor. Described as “The world’s greatest festival of ideas and discovery” this event a wide range of members of the general public and is an ideal opportunity for the ARCHER2 team to showcase the societal value of Supercomputing.
- EPCC will contribute to the full Women in HPC programme for SC’23 including a workshop on diversifying the HPC community, the annual Diversity Day (14th November), and a networking reception.
- It is expected that a range of researchers will apply for the present eCSE call to take advantage of the increase in scope to allow software development to prepare codes for Exascale. The team plan to support researchers to submit these proposals and help successful applicants get started.
- A greater number of online and face-to-face interactive training is planned for the last quarter of 2023.
- The team has two school students joining the centre for a week of work experience, providing the opportunity to work in a software development environment.

ARCHER2 Centralised CSE Team

It has been a productive period for the CSE team, with good progress on several service-improvement activities (discussed below), high-profile involvement in key conferences and events, and valuable contributions to the overall reliability and usability of the ARCHER2 service.

In support of users on the main system, the CSE team has helped to diagnose and resolve two problems that have emerged following on from the ARCHER2 software upgrade in June:

- CSE provided significant help to HPE towards diagnosing and resolving a memory leak that caused degraded job reliability and node availability issues during July.
- Further, the CSE team has supported users and contributed to the diagnosis and testing of a bug in the Lustre software which has been impacting various user groups to varying degrees. At the time of writing, a patched version of Lustre has been rolled out to address the issue, and the CSE team will monitor the situation over the coming weeks to ensure all is in order.

CSE Conferences, Workshops and Events

A key conference, during the period, was the annual RSE Conference which, this year, was held in Swansea, during 5th–7th September (with satellite event on 8th September, focused on training). The ARCHER2 service was presented to the RSE community via our booth, with the science delivered from the service on show via our series of postcards. The CSE team was very visible at the conference with a mix of presentations, discussion group and organisational contributions:

- Juan Herrera, Dimitrios Theodorakis, Colin Sauze, Joao Morado, Ed Hone, Iva Kavcic, *Modern training for modern FORTRAN*
- A. Turner, M. Weinzierl, E. Hone, *RSE HPC Community Birds of a Feather Session*
- A. Turner, C. Wood, J. Cohen, *Helping Support Reproducibility in Research: An Introduction to Containers for RSEs*
- A. Turner, T. Koskela, Ilektra Christidi, Mose Giordano, *Using ReFrame for reproducible and portable performance benchmarking*
- Juan Herrera, *ARCHER2 training design and delivery* (RSECon Satellite Training Event)

Andy Turner also attended the RSE Leaders Group meeting which preceded the RSE conference on 4th September. The CSE team were represented, at the ARCHER2 User Advisory Group meeting (online, 25th September) by Alan Simpson, Lorna Smith, and Andy Turner.

Other presentations and webinars of note during the period include:

- A. Turner, A. Jackson, A. Simpson, *Emissions and energy efficiency on large-scale high performance computing facilities: ARCHER2 UK national supercomputing service case study*, presented at the World Symposium on Climate Change Policies: Science and Technology in Support of Policy Making in Climate Change Mitigation and Adaptation, 19 Sep 2023, Online.
- A. Turner, *ARCHER2 environmental sustainability and top tips for using HPC*, University of Nottingham HPC User Conference, 20 Sep 2023, Nottingham, UK

Looking to the coming period, CSE has some significant contributions to the Supercomputing 2023 conference, which is being held in Denver, USA, during 12th–17th November.

- The work by James Richings and Oliver Brown, reported previously, towards a record-scale simulation of a quantum computer has been accepted as a workshop paper entitled *Energy Efficiency of Quantum Statevector Simulation at Scale*.
- Andy Turner, Adrian Jackson and Alan Simpson have had a paper on ARCHER2 accepted for the SC23 Sustainable Computing workshop.
- Weronika Filinger and Eleanor Broadway are also involved in the organisation of a Birds of a Feather session on outreach - "Designing HPC Outreach Activities" (<https://sc23.supercomputing.org/presentation/?id=bof135&sess=sess406>). This session, led by Weronika, focuses on aiding the community in effective outreach that is appropriate for different

age groups, scientific backgrounds, or venues. The goal of this session is to design several new activities that the community would be able to develop over the coming year.

- Weronika is also leading the preparation of two other Birds of a Feather sessions:
 - Updates from the HPC Certification Forum - <https://sc23.supercomputing.org/presentation/?id=bof143&sess=sess401>
 - Pathfinding in HPC Education and Training - <https://sc23.supercomputing.org/presentation/?id=bof140&sess=sess410>

Finally, Weronika is the Proceedings Chair for the **IEEE eScience'23 conference**, being held in Limassol, Cyprus, during 9th–13th October (<https://www.escience-conference.org/2023/>). The proceedings process, including 99 conference submissions, has been completed and the proceedings have been posted to the IEEE Xplore Digital Library.

Continual Service Improvement (CSI) Projects

Documentation enhancements to enable better utilisation during DCN upgrade

To enable users to continue to use ARCHER2 during the week beginning 18th September 2023, when the data centre network at the ACF was upgraded, the CSE team provided additional documentation on using job chaining and creating workflows involving multiple independent jobs. The aim was to allow users to confidently submit larger workloads to the ARCHER2 batch system in advance of the outage, to maintain good utilisation throughout the outage. The material has been included in the online documentation (<https://docs.archer2.ac.uk/user-guide/scheduler/#expressing-dependencies-between-jobs>). It has longer term value and is likely to be used to support advanced training courses and usage models.

ARCHER2 File System Performance Monitoring

In the previous period, we reported on work to extend the routine testing of ARCHER2, to incorporate (work) file-system performance monitoring using BenchIO. At the time, we noted a valuable next step to incorporate a measurement of metadata services (IOPS) into the monitoring and, to this end, the CSE has added IO500 measures to the Reframe suite, complementing the bandwidth data gathered from BenchIO.

The new measures are already producing value: they have been used to help investigate an I/O problem seen on the Lustre file system since the upgrade, and to check baseline performance of each component filesystem of /work.

ARCHER2 Benchmark Suite

The CSE team has begun work to revise the application benchmark suite, previously developed during the ARCHER service* to reflect topical requirements of the UK HPC community and likely future ambitions over the next few years. The scope of the suite has been expanded to cover both traditional HPC simulation requirements as well as newer workflows based on machine-learning techniques, and to accommodate use of both CPU and accelerator-based platforms. Furthermore, measurements will tackle energy use as well as performance (i.e., time to solution).

Work is progressing well and it is hoped that an outline of the required suite will be ready for internal review within CSE by the end of October, and that a first public version should be available by the end of the year.

*http://archer.ac.uk/documentation/white-papers/benchmarks/UK_National_HPC_Benchmarks.pdf

ARCHER2 Performance Report

This is the performance report for the ARCHER2 CSE Service for the Reporting Periods from July 2023 – September 2023.

The metrics were specified by UKRI 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 OY4: Service Threshold: >4.1 WD; Operating Service Level: >1.1 WD, ≤2.1 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 OY4: Service Threshold: >25.5 Working Days (WD); Operating Service Level: >10.5 WD, ≤15.5 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 OY4: Service Threshold: >56 Working Days (WD); Operating Service Level: >26 WD, ≤36 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 OY4: 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	July 2023		Aug 2023		Sept 2023		Q3 2023	
	Perf	Points	Perf	Points	Perf	Points	Perf	Points
ARCHER2_CSE_Level1 (MTR)	100%	0	100%	0	100%	0	100%	0
ARCHER2_CSE_Level2 (MTR)	100%	0	100%	0	100%	0	100%	0
ARCHER2_CSE_Level3 (MTR)	100%	0	100%	0	100%	0	100%	0
ARCHER2_CSE_TA (%)	100%	0	100%	0	100%	0	100%	0
Initial Response to Queries (%)	100%	0	100%	0	100%	0	100%	0
Query User Satisfaction (%)	100%	0	97.3%	0	100%	0	99.1%	0
Training Satisfaction (%)	100%	0	100%	0	100%	0	100%	0
Total		-11		-11		-11		-33

112 query feedback responses were received on query resolution in the Reporting Period. 99.1% of responses had a score of “Good”, “Very Good” or “Excellent”.

ARCHER2 CSE Queries

This section provides details on ARCHER2 CSE queries during the Reporting Periods from July 2023 – September 2023.

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.

July 2023					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	29	26	3	3	0.4h
ARCHER2_CSE_Level2	90	111	21	13	0.3h
ARCHER2_CSE_Level3	7	1	6	24	0.1h
ARCHER2_CSE_TA	6	6	4	13	0.1h
Aug 2023					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	56	55	4	3	0.4h
ARCHER2_CSE_Level2	92	89	24	12	0.3h
ARCHER2_CSE_Level3	1	1	6	45	0.2h
ARCHER2_CSE_TA	13	15	2	13	0.4h
Sept 2023					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	41	44	1	3	0.3h
ARCHER2_CSE_Level2	54	44	34	15	0.3h
ARCHER2_CSE_Level3	1	1	6	22	0.1h
ARCHER2_CSE_TA	11	13	0	11	0.5h
Q3 2023					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	126	125	1	3	0.4h
ARCHER2_CSE_Level2	236	244	34	13	0.3h
ARCHER2_CSE_Level3	9	3	6	30	0.1h
ARCHER2_CSE_TA	30	34	0	12	0.4h

CSE Query Categories

A total of 406 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	125	30.8%
ARCHER2_CSE_Level2	3rd party software	61	15.0%
	Batch system and queues	36	8.9%
	Software errors	36	8.9%
	Software installation	23	5.7%
	Compilers and system software	18	4.4%
	Courses	13	3.2%
	Login, passwords and ssh	12	3.0%
	Storage and compute resources	12	3.0%
	Access to services	10	2.5%
	Porting, performance and scaling	10	2.5%
	Hardware issue	5	1.2%
	eCSE applications/calls	5	1.2%
	Data transfer	1	0.2%
	Other: Queries which do not fit within other categories	1	0.2%
	Website and documentation	1	0.2%
ARCHER2_CSE_Level3	Software errors	3	0.7%
ARCHER2_CSE_TA	UKRI Grant	13	3.2%
	Pump-priming	8	2.0%
	Fellowship	8	2.0%
	EPSRC Responsive Mode	4	1.0%
	Top up	1	0.2%
Total		406	100.0%

ARCHER2 Training


As part of ARCHER2, the service has been developing and delivering a training programme for the ARCHER2 user community. During the third quarter of 2023, the CSE service has provided a total of 8.5 days of training, scheduled as follows:

Dates	Course	Location	Days	Attend
5 Jul	Python overhead: Reducing the burden	Online	0.5	32
19 Jul	Coupling LAMMPS and OpenFOAM for Multi-Scale Models	Online	0.5	48
26 Jul	Slurm: Scheduling jobs on ARCHER2	Online	0.5	43
22-23 Aug	Parallel Performance Analysis using Scalasca	Oxford	2	16
6-7 Sep	Advanced MPI	Online	2	11
21 Sep	Reproducible research using containers – Introduction to Singularity/Apptainer	Nottingham	1	9
28-29 Sep	Shared memory programming with OpenMP	London	2	12

The three webinars delivered in July were well attended (between 30-50 attendees).

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 4.2, i.e., better than “Very Good”. Users provided 15 responses, a response rate of 42%.



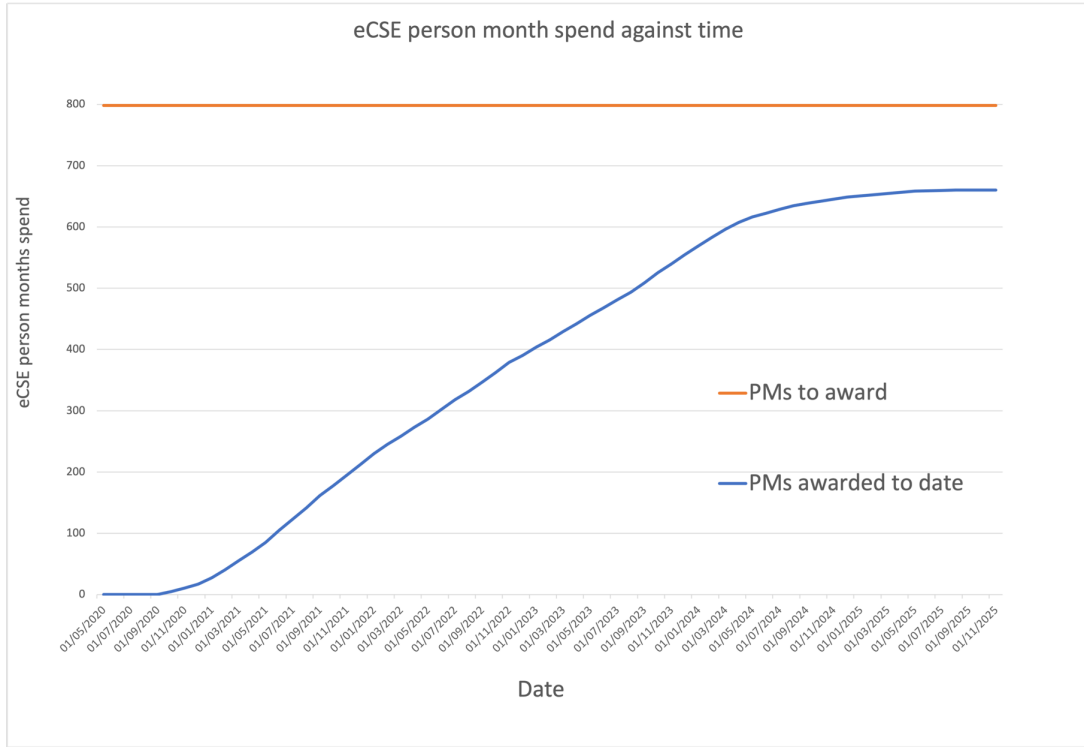
ARCHER2 Embedded CSE Programme (eCSE)

The eleventh eCSE call (eCSE11) opened on 12 September 2023 with the scope of the call widened to allow proposals for projects that will prepare software for a future Exascale system. This includes improvements to deal with issues that may only start to emerge at the Exascale level and may involve work on parallelisation/scalability, I/O, resiliency and may also include GPU development given a future Exascale system is likely to include a large GPU component. Projects involving GPU development can get access to the new GPU development cluster.

- eCSE calls 1-10
 - The eCSE programme has so far awarded 660 person months across 10 calls. The funding was awarded to 69 projects having received 116 proposals.
- For call 1 and from call 5 onwards, only proposals with software within the EPSRC remit have been eligible.
- The eCSE call 11 (eCSE11) opened on 12 September 2023, closing for technical evaluations on 10 October 2023 with final proposals due on 31 October 2023. As noted above, the scope of this call includes the development of software preparing for a future Exascale system.

eCSE call	Call Dates	# Technical Evaluations Received	# Proposals Received (EPSRC,NERC)	# PM requested (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	11 (7,4)	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	19/04/22 - 14/06/22	13	10 (10,0)	77 (77,0)	7 (7,0)	55 (55,0)
eCSE08	06/09/22- 25/10/22	17	12(12,0)	144 (144,0)	7 (7,0)	80 (80,0)
eCSE09	06/12/22- 14/03/23	12	12(12,0)	146 (146,0)	6 (6,0)	67 (67,0)
eCSE010	18/04/23 - 13/06/23	5	5 (5,0)	59 (59,0)	4 (4,0)	44 (44,0)
Total		130	116 (105,11)	1159 (1082,77)	69 (61,8)	660 (598,62)

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) including the recent extension.



ARCHER2 Community Engagement, Outreach, Collaboration and Impact

Benefits Realisation

Benefits realisation data is now available as a report from the SAFE. This is provided to UKRI on a 6-monthly basis but can additionally be run for any desired period of the service.

Blogs

Eight blogs have been published this quarter, ranging from details of our visit to the Big Bang Fair in Birmingham to do outreach activities for school children, through to articles highlighting some of the science run on ARCHER2 by our users.

Community and Outreach Activities

EPCC had a booth at the Doors Open Day at the Playfair Library in Edinburgh on the 23rd of September. Lorna Smith and Mario Antonioletti had a booth within the library, providing an overview of the ARCHER2 service and a hands-on science activity looking at sorting. The postcard series proved an excellent way to explain the science being carried out on ARCHER2 and to showcase the value of the service to the community. Over 600 members of the public visited the Playfair during the day.

Diversity and Inclusivity

This year the International HPC Summer School took place from 9 to 14 July, 2023 in Atlanta, Georgia, US, at the campus of the Georgia Institute of Technology. For the first time, the event has been expanded to include participants from Australia and observers from South Africa. Also for the first time, EPCC funded the attendance of 5 UK students (including two online MSc students). One third of the attendees were first generation college/university students. The gender ratio was slightly improved from last year for both participants and instructors. Overall, the event was very successful, with the mentoring sessions consistently receiving the highest ratings. More about this year's event can be found at <https://ss23.ihpcss.org/> and in the next EPCC News.

Weronika Filinger was also the co-chair of the Workforce Development, Training, Diversity and Education track at ACM PEARC'23 conference, held in Portland, Oregon. The role involved coordinating the review of the submitted papers, preparing the track programme and chairing the sessions. The most popular sessions had over 120 people in attendance. The conference proceedings are available at <https://dl.acm.org/doi/proceedings/10.1145/3569951>.

EPCC contributed to the preparation of a full Women in HPC programme for SC'23 (<https://womeninhpc.org/events/sc-2023>), including a workshop on diversifying the HPC community, the annual Diversity Day (Tue 14th November), and a networking reception.

Eleanor Broadway is serving as the WHPC SC'23 Submissions co-chair which includes writing and advertising a call for submissions, recruiting reviewers and conducting a peer-review process and selecting the final candidates for the WHPC workshop lightning talks. This year's WHPC SC workshop has an increased emphasis on diversity and inclusion of both women and men from under-represented groups.

The ACM SIGHPC Education Chapter is leading the effort focused on making the HPC teaching and training resources more FAIR – findable, accessible, interoperable and reusable. This includes creating a set of metadata and taxonomies that could be used to categorise training materials and ultimately create a single point of access/search.

Looking to 2024, the EPCC chapter of Women in HPC has begun planning for WHPC event to be co-located with the ARCHER2 Science Days, early in 2024. Women in HPC has also begun discussions with the International Supercomputing organisers about becoming a full partner in the conference series. The goal is to enhance the DEI efforts of the ISC conference series and to feature contributions from women working in HPC throughout the conference programme.

Quality Management, Information Security and Business Continuity

EPCC passed our ISO 9001 (quality service delivery), ISO 27001 (information security) and ISO 22301 (business continuity and disaster recovery) annual external audits with flying colours. Certification audits are an annual event, with a full re-certification every 3 years. EPCC puts in a lot of work ensuring that best practice is applied and a focus is kept on providing the highest level of service to all our users. The CSE team was heavily involved throughout the week, being audited on all aspects of the CSE service.