



# ARCHER2

## CSE Quarterly Report

July – September 2021

EPCC

The University of Edinburgh



## Document Information and Version History

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

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2021-09-29	Initial draft	Lorna Smith
0.2	2021-10-04	Added ISO, feedback and impact data	Anne Whiting
0.3	2021-10-04	Add training	Juan Rodriguez Herrera
0.4	2021-10-05	Added eCSE	Chris Johnson
0.5	2021-10-06	First full draft	Lorna Smith
1.0	2021-10-13	Version for UKRI	Alan Simpson, Lorna Smith

## 1. ARCHER 2 CSE

This is the ARCHER2 CSE report for the period July - September 2021.

### CSE Highlights

- Weronika Filinger (CSE) is the main organiser of the SC21 student program and has put together a program that would illustrate how diverse HPC is, both as a field and as a community. The aim is to show students that they can have a successful career in HPC regardless of where they live, what they look like or what their scientific background is.
- The ARCHER2 Training service was “Highly commended” in the “Practices” category of the hidden REF award ceremony, held on 2<sup>nd</sup> of September. This recognises the role training plays in facilitating high-quality research on ARCHER2.
- The CSE and SP teams produce a proposal for the initial scheduler configuration on the ARCHER2 system, based on an analysis of user requirements. This proposal was shared with expert users with their feedback showing that this met user’s needs while also providing potential improvements for further investigation.
- Four CSI projects are currently underway, looking at good practice in running containerised HPC workflows, understanding the performance of the AMD optimised compute libraries, benchmarking and optimising the performance of NEMO and NAMD, and profiling CP2K. All are designed to support and enhance the user experience.
- We have continued to deliver a high-quality training programme, with a virtual tutorial and six courses being delivered this quarter, constituting a total of 12 training days.
- The Firedrake development team, who are part of two eCSE projects, delivered one of the training courses. This allowed the user community to benefit from the eCSE work and allowed the CSE team to build closer links with the community.
- The fifth eCSE call has opened alongside an Early Career eCSE Observers call. This call gives researchers in the early stages of their careers an opportunity to attend a panel meeting to give them an insight into this competitive selection process, to assist them in the preparation of future funding proposals.
- We passed our annual external ISO 9001 Quality and 27001 Information Security audits and continue to put the highest importance on delivering the ARCHER2 service to meet the needs of our users and to ensure that all the user data entrusted to us is managed securely.
- 7 blog articles have been published this quarter highlighting the technical features and notable events happening with the service. Highlights include a series of blogs on the use of containers with ARCHER2.
- A total of 119 pieces of query feedback have been received for the CSE service from a total of 495 queries handled. All feedback was rated good and above, with the majority excellent.
- A donation of £1 per user feedback response has been made to our selected charity, Save the Children. This quarter a total donation of £228 has been made, with £119 from the CSE Query feedback.

## Forward Look

- During this quarter the full service will become available to the user community and the CSE team's main focus will be on supporting users to transition from the 4 cabinet to the full system successfully.
- The CSE team has just received access to the full cabinet system and work will focus on ensuring everything is in place to meet user needs, such as documentation, tools and support mechanisms.
- CSE are working with HPE around the evaluation of containerised versions of the HPE Cray Programming Environment. We plan to look at a range of use cases, including the use of the containerised PE in CI workflows and the use of the containerised PE to allow users to more easily access different PE versions.
- We will schedule several webinars about the ARCHER2 full system to ensure as smooth a transition as possible from the ARCHER2 4-cabinet system to the ARCHER2 23-cabinet system for users.
- We are in the process of updating training material to reflect the move to the full 23-cabinet system, for example to include the latest ARCHER2 full-system architecture and module environment.
- We are looking forward to welcoming Early Career Observers to our eCSE panel meetings. We will be inviting successful applicants to the upcoming 5<sup>th</sup> eCSE panel meeting and to subsequent eCSE panel meetings.
- With the COP26 UN Climate Conference happening in Glasgow in November 2021, several ARCHER2 related articles have been submitted for exhibition. We intend to publish these as a case study.
- The SC21 student program will run in November and, as noted above, will run virtually with a program that will illustrate how diverse HPC is, both as a field and as a community.

## 2. ARCHER2 Centralised CSE Team

During the reporting period the CSE team has worked to support the UK scientific community on the 4-cabinet ARCHER2 service, while also contributing to the implementation of the main (that is, 23-cabinet) ARCHER2 service.

As part of preparations for the main service, the CSE team worked with SP to produce a proposal for the initial scheduler configuration on the ARCHER2 system, based on an analysis of user requirements. This proposal was shared with expert users (the nominated early-access users from the Consortia) for comment and feedback. All groups confirmed that the scheduler proposal met their needs for the start of service and provided useful feedback for improvements to the scheduler that we plan to investigate, on the ARCHER2 Test and Development System, once the main service is in production.

The CSE team (along with SP) has also held regular meetings with the NCAS Computational Modelling Service and with JASMIN, to better understand data transfer requirements between ARCHER2 and JASMIN, and to collaborate more effectively on putting in place the right tools and services to support users moving data between the two services.

### CSI Projects

#### Good Practice for Running Containerised, HPC Workflows

Michael Bareford, of the CSE team, has been developing good practice for running containerised, HPC workflows on ARCHER2. Using containerised applications helps with reproducibility of results and allows users to easily run applications that are not centrally installed by the CSE team.

The work has identified an approach, using the Singularity platform, which allows users to build portable containers that can be run, with comparable performance, to a native application. This has been validated on realistic ARCHER2 job sizes (that is, thousands of cores). Furthermore, the approach can easily be carried over to other UK HPC services, and has been demonstrated for the Edinburgh-based Tier 2 (Cirrus) service.

The work is to be presented in an ARCHER2 webinar (in collaboration with the CompBioMed project) in early October, and is being written up for the ARCHER2 documentation site.

#### Understanding the Performance of the AMD Optimised Compute Libraries

The CSE service have been working with the AMD HPC Center of Excellence to understand the performance of the AMD optimised compute libraries for key applications on ARCHER2 and, where relevant, identify opportunities for performance improvements. Findings from this work will be incorporated into service documentation and training materials.

#### Benchmarking and Optimising the Performance of the Centrally Supported Codes NAMD and NEMO

In collaboration with PRACE, Eleanor Broadway is benchmarking and optimising the performance of the centrally supported codes NAMD and NEMO, on ARCHER2 (4-cabinet). The project is developing best practice for ARCHER2 users on the optimal running conditions and performance optimisations for the two applications.

We have found NAMD to require very little tuning to achieve expected performance on ARCHER2, using test cases of 8-, 28- and 210-million atoms, which all scale up to 16,384 cores.

NEMO is the second candidate code for this project. Whilst the balance of threads and cores determines the performance of NAMD, NEMO's performance is dependent on the configuration of the external XIOS servers which manage the I/O. The placement of the XIOS servers is a key aspect, which can be used to tune the performance based on the computational architecture. With these two codes we can explore different aspects of performance tuning on ARCHER2.

We plan to continue this work on the Juiot-Curie machine, a PRACE Tier-0 system, to support researchers in efficiently porting applications from ARCHER2 to Tier-0 systems.

### Profiling the CP2K Application on ARCHER2

In collaboration with the Horizon 2020 BioExcel project, Holly Judge (CSE) has been profiling the CP2K application on ARCHER2 for QM/MM systems (looking at varying the choice of functional, QM system size, and QM cell size). Holly, working with colleagues in BioExcel, will soon begin to port some of the QM/MM subroutines to use the new CP2K-internal *grid API*, which allows (for example) graphics cards and many-core processors to be targeted.

### 3. ARCHER2 CSE Performance Report

This is the performance report for the ARCHER2 CSE Service for the Reporting Periods from July 2021 until end of September 2021.

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. *Initial MTR applicable to OY1: Service Threshold: >4.4 WD; Operating Service Level: >1.4 WD, ≤2.4 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. *Initial MTR applicable to OY1: Service Threshold: >27 Working Days (WD); Operating Service Level: >12 WD, ≤17 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. *Initial MTR applicable to OY1: Service Threshold: >59 Working Days (WD); Operating Service Level: >29 WD, ≤39 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 OY1: 8 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 2021		August 2021		September 2021		Q3 2021	
	Perf	Points	Perf	Points	Perf	Points	Perf	Points
ARCHER2_CSE_Level1 (MTR)	0.200	-2	0.200	-2	0.200	-2	0.200	-6
ARCHER2_CSE_Level2 (MTR)	0.200	-2	0.200	-2	0.200	-2	0.200	-6
ARCHER2_CSE_Level3 (MTR)	0.500	-1	0.500	-0.5	0.500	-0.5	0.500	-2
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 (%)	100%	-2	100%	-2	100%	-2	100%	-6
Training Satisfaction (%)	100%	-1	100%	-1	100%	-1	100%	-3
<b>Total</b>		-10		-9.5		-9.5		-29

118 query feedback responses were received on query resolution in the Reporting Period. 100% of responses had a score of “Good”, “Very Good” or “Excellent”.



## 4. ARCHER2 CSE Queries

This section provides details on ARCHER2 CSE queries during the Reporting Periods from July 2021 until end of September 2021.

### 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.

A total of 493 queries were resolved by the ARCHER2 CSE service in the Reporting Period.

July 2021					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	146	145	1	4	0.3 hrs
ARCHER2_CSE_Level2	89	86	26	12	0.3 hrs
ARCHER2_CSE_Level3	2	2	5	30	0.2 hrs
ARCHER2_CSE_TA	11	9	3	11	0.3 hrs
August 2021					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	73	74	0	4	0.4 hrs
ARCHER2_CSE_Level2	42	48	20	16	0.3 hrs
ARCHER2_CSE_Level3	3	2	6	35	0.1 hrs
ARCHER2_CSE_TA	8	7	4	10	0.2 hrs
September 2021					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	52	51	1	4	0.4 hrs
ARCHER2_CSE_Level2	55	58	17	15	0.4 hrs
ARCHER2_CSE_Level3	0	1	5	25	0.1 hrs
ARCHER2_CSE_TA	7	10	1	16	0.5 hrs
Q3 2021					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	271	270	1	4	0.3 hrs
ARCHER2_CSE_Level2	186	192	17	14	0.3 hrs
ARCHER2_CSE_Level3	5	5	5	31	0.1 hrs
ARCHER2_CSE_TA	26	26	1	13	0.3 hrs

## CSE Query Categories

Resolved CSE queries in the Reporting Period fell into the following categories:

Service level	Category	Number resolved	% Queries
ARCHER2_CSE_Level1	Courses	270	54.8%
ARCHER2_CSE_Level2	3rd Party Software	82	16.6%
	Batch system and queues	23	4.7%
	Compilers and system software	20	4.1%
	Disk, tapes, resources	14	2.8%
	User programs	12	2.4%
	eCSE Applications/Calls	10	2.0%
	Login, passwords and ssh	9	1.8%
	Performance and scaling	8	1.6%
	Access to HPC	4	0.8%
	Courses	4	0.8%
	Porting	3	0.6%
	Data Transfer	1	0.2%
	Network	1	0.2%
	Node Failure	1	0.2%
ARCHER2_CSE_Level3	3rd Party Software	3	0.6%
	Batch system and queues	1	0.2%
	Compilers and system software	1	0.2%
ARCHER2_CSE_TA	UKRI Grant	15	3.0%
	Pump priming	6	1.2%
	EPSRC Fellowship	2	0.4%
	Non-UKRI Grant	2	0.4%
	Director Time	1	0.2%
<b>Total</b>		<b>493</b>	<b>100%</b>

## 5. ARCHER2 Training

The service has continued to develop and deliver an online training programme for the ARCHER2 community. During the second quarter of 2021, the CSE service has provided a total of 12 days of online training, scheduled as follows:

Dates	Course	Days	Attend
14, 16 Jul 2021	Advanced MPI	2	25
20-23 Jul 2021	Modern C++ for Computational Scientists	2	24
27, 29 Jul 2021	Parallel Performance Analysis using Scalasca	2	23
28-29 Jul 2021	Reproducible computational environments using containers	2	35
23 Aug 2021	Introduction to using Firedrake on ARCHER2	0.5	36
21 Aug 2021	ARCHER2 for software package users	1	16
7-8 Sep 2021	Performance Optimisation on AMD EPYC	2	48
29 Sep 2021	First year of the ARCHER2 training service	0.5	6

The ARCHER2 training service was audited during the recent ISO:9001 and ISO:27001 audits. The audit results showed that the transition from face-to-face to online courses was successful, based on the feedback received from users. The audit report describing this as “outstanding”.

The course entitled “Introduction to using Firedrake on ARCHER2” was delivered by the Firedrake development team as part of three active ARCHER2 projects (two of them ARCHER2 eCSE projects). Firedrake is an automated system for the solution of partial differential equations using the finite element method (FEM). Firedrake uses sophisticated code generation to provide mathematicians, scientists, and engineers with a very high productivity way to create sophisticated high-performance simulations. Running this course with the developers allowed the CSE team to work closely with the developers and for the benefit of the eCSE projects to be disseminated to the user community.

Several ARCHER2 CSE staff members have attended an internal training course delivered by AMD. One of the goals of this collaboration with AMD is to deliver a webinar on tips and tricks using AMD tools on ARCHER2.

The ARCHER2 Training service was “Highly commended” in the “Practices” category in the hidden REF award ceremony, held on 2<sup>nd</sup> of September. The hidden REF is a competition that recognises all research outputs and every role that makes research possible.

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 4.32, i.e., better than “Very Good”. Users provided 59 responses, a response rate of 32%.



The online self-service ARCHER2 courses on OpenMP and MPI were used as training material for postgraduate students on the two workshops “International HPC Summer School” and the “High Performance Computing Autumn Academy 2021”.

## 6. ARCHER2 The Embedded CSE Programme (eCSE)

### eCSE Calls 1-4

From the first 4 calls, 318 PMs have been awarded across 35 projects. These are detailed in the table below.

From the second call onwards, projects have been awarded to develop software from the NERC remit as well as from EPSRC. Projects from the NERC remit receive 10.73% of the awarded PMs giving 514 PMs for projects within the EPSRC remit and 62 PMs for projects within the NERC remit.

All of the 62 NERC PMs have now been allocated. From the fifth call onwards eCSE funding will be open to projects within the EPSRC remit only.

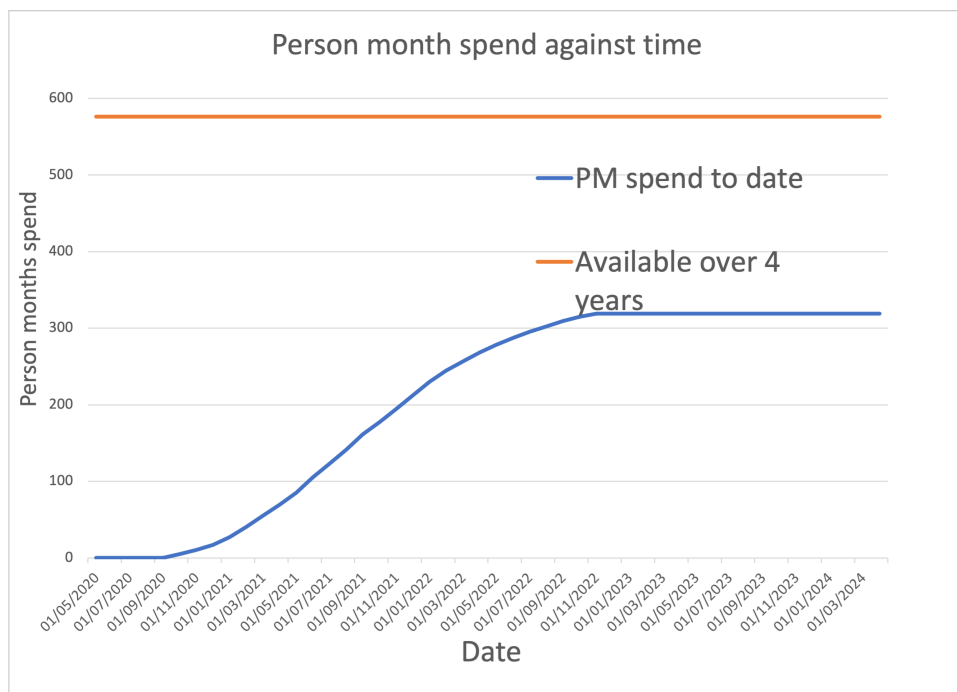
The following table gives details of the proposals received and the projects awarded across the first 4 calls. To date, 63 proposals have been received requesting a total of 478 PMs with 35 projects having been awarded a total of 318 PMs.

Some proposals fell within both EPSRC and NERC remits and were assigned a split (e.g. 80:20 EPSRC:NERC) which was agreed between the two research councils. When counting proposals or projects these have been counted according to which remit assigned the largest fraction, with the individual PMs then counted according to the agreed ratio.

One project from the eCSE01 call started considerably later than planned (July 2021) having been delayed due to a previous project being extended for COVID-19 reasons.

eCSE call	Call open date	Tech eval deadline	# Tech evals rec'd	Final call deadline	# Props rec'd (EPSRC, NERC)	# PM req'd (EPSRC, NERC)	# Props accep'd (EPSRC, NERC)	# PM Award'd (EPSRC, NERC)	# Projs started (EPSRC, NERC)	# Projs ended (EPSRC, NERC)
eCSE 01	19/05/20	16/06/20	25	07/07/20	25 (25,0)	235 (235,0)	13 (13,0)	132 (132,0)	12 (12,0)	2 (2,0)
eCSE 02	08/09/20	06/10/20	13	27/10/20	12 (9,3)	107 (87,20)	7 (4,3)	53 (33,20)	7 (4,3)	0 (0,0)
eCSE 03	08/12/20	23/02/21	15	16/03/21	14 (10,4)	136 (105,31)	8 (6,2)	73 (54,19)	3(2,1)	0 (0,0)
eCSE 04	20/04/21	18/05/21	13	08/06/21	12 (9,3)	109 (83,26)	7 (4,3)	60 (37,23)	3 (1,2)	0 (0,0)
<b>Total</b>			<b>66</b>		<b>63 (53,10)</b>	<b>587 (510,77)</b>	<b>35 (27,8)</b>	<b>318 (256,62)</b>	<b>25 (19,6)</b>	<b>0 (2,0)</b>

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).



## eCSE Call 5

The eCSE05 call opened on 7 Sep 2021 and closes for technical evaluations on 5 Oct with the final deadline for full submission being the 26 Oct 2021. An early career observers call is running alongside the call with the deadline for submissions to this call being 26 Oct 2021.

## Future eCSE Calls

Calls will be issued three times per year with a regular timetable based on similar dates to those for the calls described above. The 6<sup>th</sup> ARCHER2 eCSE call (ARCHER2 eCSE06) opens on the 7 Dec 2021 with the deadline for submitting documents for technical evaluations on 1 Feb 2021 and the final deadline for proposal submission being 16:00 on 22 Feb 2021.

## 7. ARCHER2 Community Engagement, Outreach, Collaboration and Impact

Over the last 3 months the members of ARCHER2 team have been involved in the following activities.

### Community Engagement

#### International HPC Summer School

Members of the CSE team (David Henty and Weronika Filinger) actively contributed to planning, organisation and running of the International HPC Summer school. This is a big international event that aims at equipping young scientists with skills required to effectively use HPC resources in their research. The event also fosters networking, collaboration, and community building. Students often list their mentoring experience, including near-peer mentoring, as a highlight. It was therefore an excellent opportunity for outreach and community engagement.

The event ran successfully over two weeks in 18<sup>th</sup>–30<sup>th</sup> July. The summer school is a collaboration between PRACE, XSEDE, SciNet and RIKEN, and includes students from 3 continents. This was the first time the school has been run online, which affected both the organisational aspects and the student experience. For example, participants were spread across multiple time zones forcing a split into two groups. Despite the event being held online, post-event surveys were very positive with 91% of the students saying their goals were achieved and 97% saying the code of conduct was effective at setting expectations.

The greatest challenge was creating an environment that would facilitate student interactions. The poster sessions were a good starting point (as students could share their research) but the follow-up networking was not easy virtually. Overall, the event was very successful.

#### ACM Student Research Competition @SC21

This competition showcases original research from undergraduate and graduate students, with winners selected based on poster content and presentation quality. The ARCHER2 CSE team was represented by Weronika Filinger, who was one of the reviewers.

#### Student Program @ SC21

The CSE team is also involved in the student program at SC21. Weronika Filinger (CSE) is the main organiser of the program and Ben Morse (Training) has agreed to sit on a panel dedicated to resumé/ CV writing.

The sessions for students will run for 6 days – Sunday to Friday and Weronika has put together a program that would illustrate how diverse HPC is, both as a field and as a community. For the first time ever, there will be many student programming sessions featuring so many different speakers and topics. The whole program will be virtual so the barrier to participation will be lower than ever. We want to reach as many students as possible, and show them they can have a successful career in HPC regardless of where they live, what they look like or what their scientific background is.

To make the sessions accessible to as many students as possible, activities on each day will run for around 3.5–4h. The sessions have been divided into 3 categories: Careers in HPC, Technical Skills and Non-technical skills – in total giving 16 events and 10 drop-in resumé review sessions. To show the diversity of the HPC community over 50 speakers from 6 continents were invited to participate.

## Blog Articles

We have published seven blog articles this quarter. We were able to share the arrival of the final cabinets for the full system in July, complete with photos. We have run a series of blogs on the use of containers on ARCHER2 including details of how to run a containerised GROMACS application across multiple nodes. Future blogs will highlight the full system, discuss relevant technical issues and solutions, and highlight the science being run on ARCHER2.

## ISO 9001 Quality and ISO 27001 Information Security Audits

We have once again passed our annual external ISO 9001 Quality and 27001 Information Security audits. We continue to put the highest importance on delivering the ARCHER2 service to meet the needs of our users and to ensure that all the user data entrusted to us is managed securely, not just for the external audit but throughout the year. The auditor spent a week with us looking at all aspects of what we do and how we do it. With both standards being audited at the same time, we look at everything from service metrics and user feedback to systems logs and staff employment practices.

## Presentations, Publications and Outputs

- Holly Judge, “ARCHER2 Update”, Materials Chemistry Consortium Meeting, 13 July 2021 (online)
- Andy Turner, “EPCC Facilities and HPC Tips”, University of Nottingham HPC Conference, 8 Sep 2021 (online)
- Lorna Smith “EPCC and ARCHER2”, Live at ISC21! AMD User Forum Meeting, 30<sup>th</sup> June 2021 (online)
- Andy Turner, “ARCHER2: Performance on 10,000 to (almost!) 500,000 cores”, AMD HPC User Forum Meeting, 28–29 Sep 2021 (online)