



ARCHER2 Quarterly Report

Jan – March 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	2022-03-22	Initial draft	Lorna Smith
0.2	2022-04-05	ARCHER2 CSE performance report and queries stats added	Xu Guo
0.3	2022-04-07	ISO, impact, feedback data added	Anne Whiting
0.4	2022-04-08	CSE In-depth report added	George Beckett
0.5	2022-04-08	eCSE report added	Chris Johnson
0.6	2022-04-12	Additional content and review, first full draft	Lorna Smith
0.7	2022-04-13	Reviewed	Alan Simpson
1.0	2022-04-14	Version for UKRI	Lorna Smith, Alan Simpson

ARCHER 2 Quarterly Report

This section of the report covers the period Jan 2022 – March 2022 for the ARCHER2 service.

2. ARCHER2 Executive Summary

- The CSE team has focused on supporting the user community exploit the full ARCHER2 system this quarter. Utilisation has been high on the system and the team have supported users with a wide range of queries relating to, e.g., porting, optimisation, data transfer and efficient use of the batch system.
- User feedback received on query handling has been excellent this quarter, with 100 percent of query feedback received being categorised as Good, Very Good or Excellent. 110 items of query feedback were received, with 12 Good, 17 Very Good and 81 Excellent. A donation of £305 pounds will be made to our chosen charity, Save The Children; of this £110 is from CSE query feedback.
- Users have reported an unusual number of job failures on the system. The CSE team has worked with the SP team and HPE to investigate these failures and this work is on-going. Some classes of node problems have been identified and improved node health-checking has been implemented. The CSE and SP teams continue to support users while this is on-going.
- The training team surveyed past ARCHER2 training users to understand their preferred training format. The survey showed that online training is widely preferred, with the option of having some courses face-to-face.
- Our investigations into the benchmarking, optimisation and tuning of NAMD and NEMO has been completed, and best-practice information added to the online ARCHER2 User Documentation.
- After two years of delivering courses online, the Software Carpentry workshop was delivered face-to-face at Heriot-Watt University. The course was a success, and the attendees were happy to attend the course in person.
- The training programme for the upcoming ARCHER2 training year has been approved by the ARCHER2 training panel. This programme has been tailored taking into the account the feedback provided by past attendees, the ARCHER2 Training Forum, and the ARCHER2 Training Panel.
- The sixth eCSE call closed for technical evaluations on 22nd February 2022. Five projects were awarded funding. The panel was attended by four early career observers who reported that the experience will be of benefit to them in the future.
- A total of 18 days of training, including courses and virtual tutorials, have been delivered using the full ARCHER2 system.
- As COVID restrictions have eased, the CSE team has been able to engage with the community through community events. A highlight is that the EPCC team will be well-represented at the Cray User Group meeting, in May, with five accepted papers/presentations.
- Women in HPC will be running a diversity day at ISC 22, to coincide with a half-day workshop on creating a diverse and inclusive HPC community. Our staff are involved both in organising the workshop and in running Women in HPC.

3. ARCHER2 Forward Look

- Work on ISO 22301, the business continuity standard, will continue with our stage 1 external audit from our chosen external certification body in June. The process of thinking through which elements of the services we deliver are most time critical is important for the delivery of a quality service. This should ensure minimisation of the disruption to our users by bringing back critical elements first should any disaster strike.
- A new technical mini project will commence during the next quarter, looking to develop a lightweight and flexible method 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.
- The CSE team will continue to work closely with the SP team and HPE to address the high level of job failures and to support users experiencing problems.
- The first batch of eCSE projects have finished with final reports having been received. During the next quarter, the team will concentrate on publicising highlights and analysing benefits arising from these projects.
- Now that the full ARCHER2 system is in place, we are exploring ways of maximising the appeal of the eCSE programme. We will do this for various groups, including researchers who want to scale up codes from Tier-2 systems but may require effort to evaluate and explore the scalability of their codes, and also for existing ARCHER2 users wishing to scale up to the full ARCHER2 system.
- The list of centrally installed science software is being refreshed in April, including the introduction of a new debugger/ profiler (ARM Forge), which we will evaluate the benefits of during the next 12 months.
- The CSE team has begun a new mini-project to enable fine-grain monitoring of energy utilisation of user jobs on ARCHER2.

4. ARCHER2 Centralised CSE Team

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

The CSE team has undertaken a periodic review of the centrally supported software on ARCHER2 and proposed some modest changes to maintain the relevance and usefulness of the software portfolio. The proposal was presented to the Change Advisory Board in March, has been approved, and will be implemented during April.

The CSE team has funded a twelve-month licence for the ARM Forge profiler and debugger software. This has been done to improve the quality of development tools available to users, when working at capability scale. During the licence term, the CSE team plan to monitor benefits and, if sufficiently high, to seek funding to extend the licence for the lifetime of the service. The ARM Forge software is being added to the centrally supported software list and will be available to all users later in April.

The CSE team is also expecting to make a significant contribution to the Cray User Group meeting (CUG'22) which will be held in Monterey Bay, California, in May. Five different papers/presentations have been accepted from EPCC staff. Further, Juan Rodriguez Herrera has been invited to serve on the programme committee for the conference.

CSI Projects

Benchmarking and optimising the performance of centrally supported codes NAMD and NEMO

Eleanor Broadway (through the PRACE project) has completed her investigations into the benchmarking, optimisation and tuning of NAMD and NEMO, based on experience from ARCHER2. A best-practice guide has been produced and added to the online ARCHER2 User Documentation.

Work on NAMD performance tests will continue, using the PRACE Tier-0 system, Joliot-Curie, for comparison. The Joliot-Curie system has a similar processor architecture to ARCHER2 but a different network, so this should produce some interesting results on topical interconnect technologies at capability scale.

HPE CPE Containers Customer Evaluation Programme

During the period, the CSE team (mostly Michael Bareford and Andrew Turner) has continued to participate in the HPE CPE Containers Customer Evaluation Programme. A containerised version of the Cray Programming Environment Version 22.02 has been made available to CSE, which is a useful reference when reviewing the PE provision on the ARCHER2 system. Some aspects of the project will be easier to address once the ARCHER2 Test and Development System is available.

Porting and Optimising CP2K for the ARCHER2 Main System

In a collaboration between the CSE team and the BioExcel project, Holly Judge continued her investigation into the Grid API port though, due to how the QM/MM part is implemented, she determined that using the current grid API with it is very inefficient. The conclusion from that is that the QM/MM part will need its own dedicated GPU code, though that is beyond the scope of this activity. The CP2K developers note they plan to look at this in the future.

Holly also tested a build of CP2K with the HIP framework on an AMD accelerator cloud, to which she had temporary access. Testing results were mixed (possibly because this system was in development). The build instructions will go into a best practice guide (and are expected to be a useful starting point for when GPUs are added to ARCHER2). The case is short-listed for further investigation, if CSE get access to other GPU resources via an arrangement with HPE.

Monitoring the ARCHER2 Power Management Hardware Counters

This is a new technical mini project to develop a lightweight and flexible method for gathering power performance statistics generated by parallel codes. This method should be applicable to any MPI code and should collate automatically the data gathered from multi-node runs.

The project is being undertaken jointly by CSE (Michael Bareford) and HPE Centre of Excellence (Harvey Richardson) building on previous work undertaken on ARCHER, funded by the EU FP7 CRESTA Project, and extended by Michael (within ARCHER CSE scope) to capture and report energy usage.

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.

5. ARCHER2 Performance Report

This is the performance report for the ARCHER2 CSE Service for the Reporting Periods from Jan 2022 until end of March 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. *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	January 2022		February 2022		March 2022		Q1 2022	
	Perf	Points	Perf	Points	Perf	Points	Perf	Points
ARCHER2_CSE_Level1 (MTR)	1.1W1	-2	1.1W1	-2	1.1W1	-2	1.1W1	-6
ARCHER2_CSE_Level2 (MTR)	1.3W1	-2	1.3W1	-2	1.3W1	-2	1.3W1	-6
ARCHER2_CSE_Level3 (MTR)	1.0W1	-2	1.0W1	-2	-	-	1.0W1	-4
ARCHER2_CSE_TA (%)	1.00%	-1	1.00%	-1	1.00%	-1	1.00%	-3
Initial Response to Queries (%)	1.1%	-1	1.00%	-1	1.1%	-1	1.1%	-3
Query User Satisfaction (%)	1.0%	-2	1.00%	-2	1.0%	-2	1.0%	-6
Training Satisfaction (%)	1.00%	-1	95.43%	-0.25	1.00%	-1	95.43%	-2.25
Total		-11		-10.25		-9		-30.25

106 query feedback responses were received on query resolution in the Reporting Period. 100% 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 October 2021 until end of December 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.

January 2022					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	127	127	0	3	0.2hrs
ARCHER2_CSE_Level2	77	79	24	11	0.3hrs
ARCHER2_CSE_Level3	3	3	4	54	0.3hrs
ARCHER2_CSE_TA	0	1	0	16	1.9hrs
February 2022					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	175	172	3	3	0.2hrs
ARCHER2_CSE_Level2	76	64	36	11	0.4hrs
ARCHER2_CSE_Level3	2	2	4	42	0.2hrs
ARCHER2_CSE_TA	4	2	2	17	0.2hrs
March 2022					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	160	163	0	3	0.3hrs
ARCHER2_CSE_Level2	74	78	32	12	0.3hrs
ARCHER2_CSE_Level3	0	0	4	0	-
ARCHER2_CSE_TA	4	6	0	12	0.2hrs
Q1 2022					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
ARCHER2_CSE_Level1	462	462	0	3	0.3hrs
ARCHER2_CSE_Level2	227	221	32	11	0.3hrs
ARCHER2_CSE_Level3	5	5	4	49	0.3hrs
ARCHER2_CSE_TA	8	9	0	14	0.4hrs

CSE Query Categories

A total of 697 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	462	66.3%
ARCHER2_CSE_Level2	3rd Party Software	90	12.9%
	Batch system and queues	30	4.3%
	Login, passwords and ssh	22	3.2%
	Compilers and system software	19	2.7%
	eCSE Applications/Calls	15	2.2%
	Courses	12	1.7%
	User programs	12	1.7%
	Data Transfer	7	1.0%
	Performance and scaling	6	0.9%
	Porting	3	0.4%
	Access to HPC	2	0.3%
	Disk, tapes, resources	2	0.3%
	Node Failure	1	0.1%
ARCHER2_CSE_Level3	3rd Party Software	3	0.4%
	Porting	1	0.1%
	User programs	1	0.1%
ARCHER2_CSE_TA	UKRI Grant	5	0.7%
	Pump priming	3	0.4%
	EPSRC Fellowship	1	0.1%
Total		697	100%

7. ARCHER2 Training

As part of ARCHER2, the service has been developing and delivering online a training programme for the ARCHER2 community. During the first quarter of 2022, the CSE service has provided a total of 18 days of online training, scheduled as follows:

Dates	Course	Days	Attend
11 Jan 22	Efficient Parallel IO on ARCHER2	1	15
12 Jan 22	Modelling volcanic plumes using ARCHER2	0.5	17
13 Jan 22	Integrating DEC into the ParaFEM library	0.5	30
19 Jan 22	High-perf. materials modelling using CASTEP on ARCHER2	0.5	23
19-20 Jan 22	Reproducible computational environments using containers	2	31
27-28 Jan 22	HPC Carpentry	2	21
31-1 Feb 22	Introduction to ARCHER2 for Software Developers	2	7
23-24 Feb 22	Scientific Programming with Python	2	24
28 Feb 22	ARCHER2 for Data Scientists	1	17
2 Mar 22	The Hitchhikers' Guide to ARCHER2	0.5	36
21-22 Mar 22	Software Carpentry	2	21
23-24-31 Mar 22	Message-passing programming with MPI	3	22
31 Mar 22	Advanced use of LAMMPS	1	33

The number of courses and virtual tutorials has increased this quarter with respect to past quarters due to the availability of the full ARCHER2 system, which allows delivery of further content that is tailored to the characteristics and large capacity of the full system.

A survey to evaluate the training format was circulated among past ARCHER2 training users. A total of 50 responses were analysed. The overall conclusion is that the online training format is widely preferred, with the option of having some courses face-to-face.

After two years of delivering courses online, the Software Carpentry workshop was delivered face-to-face at Heriot-Watt University, given that the restrictions are eased and the situation with respect to the pandemic is less severe than before. The course was a success, and the attendees were happy to attend the course in person.

The training programme for the upcoming ARCHER2 training year has been approved by the ARCHER2 training panel. The programme has been tailored taking into the account the feedback provided by past attendees, the ARCHER2 Training Forum, and the ARCHER2 Training Panel. It includes new software package courses to cover users' needs.

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 4.52, i.e., better than "Very Good". Users provided 60 responses, a response rate of 31%. The "very bad" feedback response was a misunderstanding from the user, who attended the second day of the course thinking it was the first one. The user claimed the course was incomplete.



8. ARCHER2 Embedded CSE Programme (eCSE)

The sixth eCSE call closed for technical evaluations on 22 February 2022 and received 7 documents for technical evaluation, with 6 final proposals submitted on the 15 March. At the panel meeting on 6th April 2022, 5 of projects were funded, awarding a total of 49 PMs. One proposal was recommended for resubmission.

The sixth eCSE panel was attended by four early career observers who reported that they felt the experience was very useful to them and will clearly be of benefit to them in the future.

Now that the full ARCHER2 system is in place, we are exploring ways of maximising the appeal of the eCSE programme. We will do this for various groups, including researchers who want to scale up codes from Tier-2 systems but may require effort to evaluate and explore the scalability of their codes, and also for existing ARCHER2 users wishing to scale up to the full ARCHER2 system.


eCSE calls 1-6

Six calls have now been completed. From these six calls, 414 PMs have been awarded across 45 projects. These are detailed in the table below.

From call 5 onwards, only proposals with software producing science within the EPSRC remit have been eligible.

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)
Total		83	78 (68,10)	773 (656,77)	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).



eCSE call 7

The seventh eCSE call opens on 19th April 2022, closes for technical evaluations on 17th May 2022 and the final deadline for proposals is 7th June 2022.

9. ARCHER2 Community Engagement, Outreach, Collaboration and Impact

Community Engagement and Impact

Eleanor Broadway has taken on the role of Volunteer Chair for the 2022 run of the Research Software Engineering conference, which will be held in Newcastle in August. Preparatory work has begun, and Eleanor is developing policy documents and role descriptions for this.

Weronika Filingier has been part of a small team from the ACM SIGHPC Education Chapter that, for over a year, has been working with relevant groups within ACM to create an award for outstanding contributions to computational science education. The award has now been approved, with a goal to recognise those who have contributed to computational science education through one or more projects or programs that have had a significant impact on the number, quality, and diversity of those applying computational science in their field(s) of study. Following on from this, Weronika has also been appointed to the SIGHPC Education Awards Committee.

David Henty has been invited to review applications from prospective European attendees for the International HPC Summer School, taking place in Athens during the 19–24 June. The event's agenda is being finalised and it has been confirmed that David will run the MPI sessions, Weronika Filingier will run the mentoring sessions, and Ludovic Capelli (EPCC Teaching Assistant) is responsible for the Programming Challenge.

The CSE team contributed to two successful Birds of a Feather proposals to ISC22. The first proposal focuses on International HPC Certification with the goal to ease the provision and uptake of training by clearly categorising, defining, and eventually assessing the skills required to efficiently use HPC resources. The second BoF is on “Software Engineering and Reuse in Modelling, Simulations and Data Analytics for Science and Engineering”, aiming to raise awareness on the importance of software engineering for computational science and engineering and its challenges.

In the scope of Women in HPC, Weronika Filingier and George Beckett are both involved in organising the “Women in HPC: creating a diverse and inclusive community” workshop at ISC'22.

The CSE team also contributed to the following community events:

- Andrew Turner participated in the 32nd Materials Chemistry Consortium Meeting, held online on 27th January.
- Weronika Filingier was invited to present and be part of a panel discussion at the Supercomputing Asia conference (NSCC, Singapore + hybrid) in the “Accelerating HPC Upskilling Without Borders” track, running on 2nd March. The presentation was on “HPC Education and Training at EPCC”, giving an overview of EPCC's education and training ecosystem – from ARCHER2 training, through the MSC programmes to other local and international engagements.
- Andrew Turner participated in a panel session on “Anticipating Researcher Needs” at the N8 Digital Research Infrastructure Retreat, in Manchester University, on 29th March.

Outreach & Engagement

Lorna Smith is organising an ARCHER2 booth at the Big Bang Fair in Birmingham in June this year, highlighting the value of Supercomputing to society and engaging young people in the benefits of a career in computational science.

Blog Posts

We have published six blogs this quarter on a range of topics:

- Disaster recovery testing to improve resilience
- Launch of the ARCHER2 Driving Test
- Software data usage on ARCHER2
- Computing Insite UK – back in person
- Investigating the performance of NAMD and NEMO on ARCHER2
- Commercial access to ARCHER2 and Cirrus

Work continues to ensure that we are collecting the correct data from the SAFE on the service to best support UKRI in demonstrating the value of investment in HPC for science for the future.

Quality Management and Information Security

Preparation work continues to develop a Business Continuity Management System on our journey to ISO 22301 certification. The process of thinking through which parts of the services we deliver are most time critical is important for the delivery of a quality service. This should ensure minimisation of the disruption to our users by bringing back critical parts first should any disaster strike. A review has been carried out on the ACF data centre and, not surprisingly, highest on the list was access to the building for all relevant staff so that they can assess the situation and start work resolving issues.

With the increased cyber threat warnings being sent out from the National Cyber Security Centre due to the war in Ukraine, we have been reviewing our information security arrangements and running a series of seminars for our staff to ensure everyone is aware of the current situation and knows how best they can help prevent any exploits of vulnerabilities in our environments. Our staff are involved in the University's discussions and preventative measures and also ensure that they are subscribed and involved with external bodies to exchange information and best practice.

We continue to work through our annual programme of internal audits for both quality service delivery and information security, highlighting good practice and looking for improvement suggestions, which we then track to make sure they are effective.