



ARCHER2

SP Quarterly Report

July – September 2024
EPCC
The University of Edinburgh



Document Information and Version History

Version:	1.0
Status	Release
Author(s):	Clair Barrass, Jo Beech-Brandt, Alan Simpson, Anne Whiting
Reviewer(s)	Alan Simpson

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	24/09/2024	Template created	Jo Beech-Brandt
0.2	02/10/2024	Added heatmap and usage graphs	Clair Barrass
0.3	02/10/2024	Added user satisfaction and ISO info	Anne Whiting
0.4	03/10/2024	Added narrative, metrics and graphs	Jo Beech-Brandt
0.5	04/10/2024	Added hosting update	Paul Clark
0.6	09/10/2024	Reviewed	Alan Simpson
1.0	/09/10/2024	Version for UKRI	Alan Simpson, Jo Beech-Brandt

1 The ARCHER2 Service

This is the report for the ARCHER2 SP Service for the Reporting Period: 1 July – 30 September 2024.

1.1 Service Highlights

- Capability Days 3 took place from 24-26 September 2024 with improvements following feedback from Capability Days 2. This included a pre-capability day reservation and a special reservation in place for NERC users.
- The latest Access to HPC Call applications are now all on boarded on ARCHER2. This call allowed us to use a new SAFE process where applicants were asked to complete their Technical Assessments directly submitting details into the SAFE. This helped to streamline the process and allowed ARCHER2 staff to complete the Technical Assessments within the SAFE also. Applicants were emailed with details of the new process and user documentation was prepared for this new process also.
- Globus Online data transfer capacity has been advertised to users, and users are now making use of it.
- Ticketing infrastructure for a new Slurm fairshare configuration, linking scheduling share to project allocation, has been prepared and testing is ongoing.
- We are working towards our next business continuity test to identify any improvements needed to make the ARCHER2 service as resilient as possible.

1.2 Forward Look

- Webinar material especially aimed at PIs and the use of SAFE to manage their projects has been produced. This webinar will be delivered during the next quarter. This is in addition to the specific documentation for PIs.
- EPCC are engaged in high-level discussions with HPE about the future of the system software. The aim is to ensure good information security throughout the service, while minimising disruption to users.
- ARCHER2 service staff will be attending SuperComputing in Atlanta and the CIUK Conference in Manchester. ARCHER2 staff will also attend the HPC-AI Advisory Council meeting and the ExCALIBUR workshop in Leicester. Plans continue for the next ARCHER2 Celebration of Science event.
- We expect to deploy the new Slurm fairshare configuration, linking scheduling share to project allocation, during Q4 of this year.
- A new training registration process will be implemented which will allow users to register for training courses directly within the SAFE. This will create a more streamlined process for users, and also ARCHER2 staff, and allow easier management of courses.

2 ARCHER2 Performance Report

This is the contractual performance report for the ARCHER2 SP Service for the Reporting Periods from 1 July 2024 until 30 September 2024.

2.1 Service Points and Service Credits

The Service Levels and Service Points for the SP service are defined by EPSRC in Schedule 2.2 of ARCHER2 SP Service Contract.

The Working Day (WD) for the ARCHER2 Service is 10 Working Hours (WH) as the Service operates from 0800-1800. The Median Time to Resolution is measured in WD.

- **Availability:** *Service Threshold: $\leq 96.5\%$; Operating Service Level: $>98.0\%$, $\leq 98.5\%$.*
- **ARCHER2_SP_Level1 (MTR):** *The Median Time to Resolution, of all SP queries falling within Level 1 resolved by the Contractor in the Reporting Period. MTR Service Threshold: >1 WD; Operating Service Level: >0.3 WD, ≤ 0.45 WD.*
- **ARCHER2_SP_Level2 (MTR):** *The Median Time to Resolution, of all SP queries falling within Level 2 resolved by the Contractor in the Reporting Period. MTR Service Threshold: >8 WD; Operating Service Level: >2 WD, ≤ 4 WD.*
- **ARCHER2_SP_Level3 (MTR):** *The Median Time to Resolution, of all SP queries falling within Level 3 resolved by the Contractor in the Reporting Period. MTR Service Threshold: >25 WD; Operating Service Level: >12 WD, ≤ 16 WD.*
- **Initial Response to Queries (%):** *The percentage of the total number of SP 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: $82.00 - 87.99\%$*

2.1.1 Service Points

Metric	Jul 2024		Aug 2024		Sep 2024		Q3 2024	
	Perf	Points	Perf	Points	Perf	Points	Perf	Points
Availability	100%	-3	99.8%	-3	100%	-3	99.9%	-9
SP_Level1 (MTR)	0.00	-2	0.00	-2	0.00	-2	0.00	-6
SP_Level2 (MTR)	0.07	-2	0.08	-2	0.09	-2	0.08	-6
SP_Level3 (MTR)	7.89	-2	0.00	-2	4.51	-2	5.41	-6
Initial Response (%)	100%	-1	100%	-1	100%	-1	100%	-3
Query Satisfaction (%)	100%	-2	100%	-2	100%	-2	100%	-6
Total		-12		-12		-12		-36

2.1.2 Service Credits

As the Total Service Points are negative (-36), no Service Credits apply in 24Q3.

2.2 SP Query Statistics

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

- **Assigned:** The number of SP queries assigned to the Contractor within each query resolution category in the Reporting Period.
- **Resolved:** The number of SP queries resolved by the Contractor within each query resolution category in the Reporting Period.
- **Backlog:** The number of SP 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 SP queries in each query resolution category.
- **First Response:** The average time taken for the Contractor to first respond to the Originator of the SP query.

July 2024					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
SP_Level1	1483	1480	3	0.06	0:00:50
SP_Level2	86	71	40	6.746	0:13:34
SP_Level3	0	2	0	15	0:29:10
August 2024					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
SP_Level1	1568	1571	0	0.092	0:00:47
SP_Level2	69	66	43	7.712	0:14:24
SP_Level3	1	0	1	0	0:00:00
September 2024					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
SP_Level1	1589	1589	0	0.086	0:00:54
SP_Level2	69	96	16	7.521	0:13:41
SP_Level3	1	1	1	18	0:04:21
Q3 2024					
Service level	Assigned	Resolved	Backlog	Correspondence	First Response
SP_Level1	4640	4640	0	0.08	0:0:51
SP_Level2	224	233	16	7.339	0:13:51
SP_Level3	2	3	1	16	0:20:54

2.3 Query Resolution

Metric	Jul 2024		Aug 2024		Sep 2024		Q3 2024	
	MTR	Resolved	MTR	Resolved	MTR	Resolved	MTR	Resolved
SP_Level1	0:00:10	1480	0:00:13	1571	0:00:30	1589	0:00:14	4640
SP_Level2	0:42:54	71	0:46:17	66	0:56:25	96	0:47:57	233
SP_Level3	78:54:30	2	0:00:00	0	45:05:46	1	54:07:11	3
Total		1553		1637		1686		4876

A total of 4876 queries were resolved by the ARCHER2 SP Service in the Reporting Period. The percentage of user queries responded to within 3 hours was 100%.

2.4 Query Feedback

During July, there were 29 feedback scores received during this period. 100% were Good, Very Good or Excellent with 76% given the highest score of Excellent.

During August there were 21 feedback scores received during this period. 100% were Good, Very Good or Excellent with 86% given the highest score of Excellent.

During September there were 23 feedback scores received during this period. 100% were Good, Very Good or Excellent with 87% given the highest score of Excellent.

£73 Donation was made to our chosen charity Save the Children with £1 donated per query feedback item received.

2.5 Maintenance and Outages


Type	Start	End	Duration	User Impact	Reason	Attributable
Partial	2024-09-30 0800	2024-09-30 1400	6 hrs	fs4 - Some data or directories on the file system may be inaccessible. Trying to access inaccessible data may cause the terminal to hang.	OSS failure and failover did not happen successfully	HPE
Partial	2024-09-25 1600	2024-09-24 1500	23 hrs	Slurm scheduler- Intermittent issues running Slurm commands	Aggressive polling of the slurm queue and user workflow	User load
Partial	2024-09-24 1530	2024-09-24 1700	1 hr 30 mins	fs1- Slow response to access data on fs1 work file system. `module` commands show slow response.	Contention for file system resources	User load
Partial	2024-07-18 20:30	2024-07-19 12:00	15 hrs 30 mins	fs3 - Issues accessing data on work (fs3) file system	Contention for file system resources	User load

The issues attributed to “user load” was due to user contention on the file system. It should be noted that two of the outages occurred during capability days when the system was being stressed. The service remained available to users but accessing data on the file system was slow at times.

3 ARCHER2 Service Statistics


3.1 Utilisation

Utilisation from 1 July – 30 September is 94% which is slightly increased from 92% the previous quarter. Utilisation for July was 89%, for August 97% and for September 90%.




The utilisation by the Research Councils, relative to their respective allocations, is presented below. This bar chart shows the usage of ARCHER2 by the two Research Councils presented as a percentage of the total Research Council allocation on ARCHER2. It can be seen that EPSRC exceeded their target this quarter with their usage being at 73% (against their target of 66.8%). It should be noted that the proportion of EPSRC's uncharged utilisation increased this quarter and is 17%.

NERC also exceeded their target with utilisation being 21% (against their target of 18.2%) which was an increase from 15% in the previous quarter.




The stacked graph below shows the trend of charge and uncharged utilisation since the start of the service. There is a trend that the NERC usage is increasing.




3.2 Scheduling Coefficient Matrix

The colour in the matrix indicates the value of the Scheduling Coefficient. This is defined as the ratio of runtime to runtime plus wait time. Hence, a value of 1 (green) indicates that a job ran with no time waiting in the queue, a value of 0.5 (pale yellow) indicates a job queued for the same amount of time that it ran, and anything below 0.5 (orange to red) indicates that a job queued for longer than it ran.



The usage heatmap below provides an overview of the usage on ARCHER2 over the quarter for different job sizes/lengths. The colour in the heatmap indicates the number of CUs expended for each class, and the number in the box is the number of jobs of that class.

It should be noted that there was an increase in the number of larger sized jobs during this quarter as users were encouraged to submit larger jobs during the data centre network maintenance session.



Appendix: Critical Success Factors

1. Context

EPCC have been asked by UKRI to provide quarterly data for a number of critical success factors:

- CSF04 Implementation of environmentally considerate energy policies
- CSF07 Deliver and maintain a reliable data I/O function
- CSF08 Be cost-effective, cost-efficient and drive towards lowering of operational costs

In the sections below, please find the relevant metrics and data.

2. CSF04 Implementation of environmentally considerate energy policies

Implementation of environmentally considerate energy policies with a drive to reducing costs and environmental impacts.

All electricity provided to the ACF and ARCHER2 is on a 100% green, renewable energy tariff.

Environmentally considerate policies: 3

Since the start of full Service, EPCC have worked on implementing the following policies:

- Move from High Performance Mode to Low Power Mode: reduced average power draw from 3.2 MW to 2.9 MW (9%) with negligible input on performance [May 2022]
- Reduced default processor frequency: further reduced average power to around 2.5 MW (19%) [December 2022]
- Increase in coolant temperatures: this will result in an increase in passive cooling (“free cooling”) [ongoing]

Power Usage

	4Q 21*	1Q 22	2Q 22	3Q 22	4Q 22	1Q 23	2Q 23	3Q 23	4Q 23	1Q 24	2Q 24	3Q 24
Average Power	3.31	3.16	3.15	2.86	2.90	2.51	2.56	2.46	2.53	2.58	2.54	2.64

* Partial

So far, the average power draw has been reduced by around 0.7MW (21%) which will reduce electricity usage by up to 6M kWh per annum, significantly reducing annual running costs.

3. CSF07 Deliver and maintain a reliable data I/O function


The compute resource will deliver and maintain an efficient, effective and reliable data I/O function which meets the requirements of users and their software. It will evolve and expand to accommodate new software or hardware architectures as required by the Service or its user base.

Data Transferred

EPCC monitor the data transfer rates in and out of the ARCHER2 system. Based on this, we now estimate the total amount of data transferred on and off ARCHER2 each Quarter.

Data Transferred...	4Q 21*	1Q 22	2Q 22	3Q 22	4Q 22	1Q 23	2Q 23	3Q 23	4Q 23	1Q 24	2Q 24	3Q 24
...to ARCHER2 (TB)	534	163	68	220	44	67	42	65	99	108	93	98
...from ARCHER2 (TB)	236	582	667	822	834	1231	1022	1472	1771	2056	2443	3956

* Partial



The trend of more data being moved off each quarter continues. Indeed, there was more data transferred off ARCHER2 in the last quarter as there was in the first year of the ARCHER2 service.

Parallel IO Write Performance

We regularly monitor the parallel write performance between the compute nodes and the parallel Lustre (/work) file systems. We use the benchio synthetic IO benchmark application (<https://github.com/davidhenty/benchio>) and report the MPI-IO write performance with the following settings:


- Global data structure of 20483: writes a single file of 65,536 MiB (64 GiB).
- Uses 16 compute nodes and 128 MPI processes per node.
- Uses UCX as the MPI transport protocol.
- Sets the following environment variables:
 - FI_OFI_RXM_SAR_LIMIT=64K

- MPICH_MPIO_HINTS="*:cray_cb_write_lock_mode=2,*:cray_cb_nodes_multiplier=4

These settings have been found to maximise the IO performance for parallel writes using MPI-IO on the ARCHER2 file systems. Writes using the default settings on ARCHER2 typically have median write values 2-3 GiB/s lower than the optimised values.

Original reporting of this data (Q1 and Q2 2023) used the means from a small number of runs on the HDD-based Lustre file systems. From Q3 2023 onwards we have been monitoring performance regularly on both HDD and NVMe-based Lustre file systems throughout the quarter and report median (Q2) and lower (Q1) and upper quartile (Q3) performance and provide boxplots illustrating the performance variation. (On the boxplots, the green triangles mark the mean value and the whiskers extend to the last datapoint within the range 1.5 x IQR.)

Benchio MPI-IO medium (GiB/s)	1Q23	2Q23	3Q23	4Q23	1Q24	2Q24	3Q24
a2fs-work1	8.2	7.6±0.5	10.5 (8.8:11.8)	10.9 (8.3:12.5)		10.1 (7.0:11.8)	9.7 (6.7:11.9)
a2fs-work2	8.5	7.3±0.6	10.4 (7.2:12.4)	10.4 (7.7:13.0)		11.1 (8.0:12.5)	11.1 (8.1:13.1)
a2fs-work3	8.3	9.6±0.7	10.0 (8.2:11.6)	10.7 (8.1:11.9)		9.6 (8.4:11.8)	9.6 (7.5:11.8)
a2fs-work4				9.7 (9.1:10.2)		10.0 (9.2:10.8)	10.6 (9.4:11.6)
a2fs-nvme			10.1 (9.6:11.5)	10.1 (9.5:12.4)		11.1 (10.5:12.4)	11.1 (12.7:11.6)



4. CSF08 Be cost-effective, cost-efficient and drive towards lowering of operational costs

The Service shall be cost-effective and cost-efficient across its elements during its lifetime and drive towards lowering of operational costs by seeking efficiencies in delivery such that TCO presents an acceptable and cost-effective solution for the public. The Service will monitor and report its Power Usage Effectiveness (PUE) and strive to make efficiency savings where possible.

Relative Research Output

Measure	11/2021 – 5/2022	5/2022 – 12/2022	1Q 23	2Q 23	3Q 23	4Q 23	1Q 24	2Q 24	3Q 24
Relative Research Output per kWh	100	109	115	115	115	115	115	115	115

We define the initial measure of research output per kWh on ARCHER2 to be 100, and then estimate how this has changed with the introduction of the various environmentally considerate policies discussed under CSF04. This is estimated using applications benchmarks similar to those defined by UKRI for the procurement.

Energy Used per CU Delivered

	4Q21*	1Q22	2Q22	3Q22	4Q22	1Q23	2Q23	3Q23	4Q23	1Q24	2Q24	3Q24*
Energy per CU (kWh)	0.719	0.713	0.728	0.715	0.650	0.545	0.669	0.590	0.568	0.582	0.585	0.601

*partial

Energy Cost per CU Delivered

	4Q21*	1Q22	2Q22	3Q22	4Q22	1Q23	2Q23	3Q23	4Q23	1Q24	2Q24	3Q24*
Cost per CU (£)	£0.089	£0.090	£0.098	£0.096	£0.088	£0.074	£0.162	£0.143	£0.136	£0.140	£0.160	£0.164

*partial

The two tables above are calculated using the total CUs delivered by ARCHER2, the total kWh of electricity consumed, and the unit cost for kWh. The increase in “Energy Cost per CU Delivered” from 2Q23 is caused by a significant increase in the unit cost of electricity from April 2023. For 2Q23, there is also an impact on the “Energy Used per CU Delivered” from the major software upgrade that took 3 weeks. There was also an additional increase in the unit cost of electricity from April 2024.